Efficiency and Optimality of Largest Deficit First Prioritization: Resource Allocation for Real-Time Applications

YUHUAN DU and GUSTAVO DE VECIANA, The University of Texas at Austin

An increasing number of real-time applications with compute and/or communication deadlines are being supported on shared infrastructure. Such applications can often tolerate occasional deadline violations without substantially impacting their Quality of Service (QoS). A fundamental problem in such systems is deciding how to allocate shared resources so as to meet applications' QoS requirements. A simple framework to address this problem is to, (1) dynamically prioritize users as a possibly complex function of their deficits (difference of achieved vs required QoS), and (2) allocate resources so to expedite users with higher priority. This paper focuses on a general class of systems using such priority-based resource allocation. We first characterize the set of feasible QoS requirements and show the optimality of max weight-like prioritization. We then consider simple weighted Largest Deficit First (w-LDF) prioritization policies, where users with higher weighted QoS deficits are given higher priority. The paper gives an inner bound for the feasible set under w-LDF policies, and, under an additional monotonicity assumption, characterizes its geometry leading to a sufficient condition for optimality. Additional insights on the efficiency ratio of w-LDF policies, the optimality of hierarchical-LDF and characterization of clustering of failures are also discussed.

Additional Key Words and Phrases: Soft real-time applications, cloud-computing, largest deficit first prioritization, feasibility region, feasibility optimal, geometry of inner bound, class-based hierarchical prioritization

arXiv:1601.06331v2 [cs.NI] 7 Mar 2016

1. INTRODUCTION

A growing number of real-time applications with compute and/or communication deadlines are being moved onto shared infrastructure, e.g., ranging from embedded systems to efficient cloud infrastructure. Such applications include control, multimedia processing, and/or machine learning components associated with enabling various types of user services as well as wireless, intelligent transportation and energy systems. In many cases such applications can tolerate occasional deadline violations, i.e., have soft constraints, without impacting the application Quality of Service (QoS). For example, applications with feedback can quickly compensate for errors, or humans may tolerate occasional failures in video processing since they can be partially concealed, or wireless base stations can tolerate occasional frame losses, since these can be retransmitted. More generally real-time applications' long-term QoS may depend in a complex manner on what was accomplished on time, e.g., partial completion of a set of tasks, or notions of video quality.

Enabling efficient sharing of compute/communication resources is a challenging problem. On the one hand, even for a single resource, tying the sharing model, e.g., round robin, priority schemes, to QoS metrics is generally hard due to the uncertainty in applications' workloads and possible variations in processing speeds. On the other hand, today's applications leverage complex networks of heterogeneous compute/communication resources, e.g., multi-core computers, embedded network system, or combinations of computation on mobile devices and the cloud. Consider the exam-

This research was supported by Huawei Technologies Co. Ltd.
A conference version of this paper has been accepted to INFOCOM 2016.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© YYYY ACM. 2376-3639/YYYY/01-ARTA $15.00
DOI: http://dx.doi.org/10.1145/0000000.0000000

ple in Figure 1. User 1 periodically generates a task that needs to be processed sequentially on Resources A, B, C, D in each period while User 2 generates tasks to be processed on Resource B then C. How should one go about designing resource sharing policies across multiple heterogeneous resources, where parallelism, task preemption and migration are allowed? Furthermore, how can one address heterogeneous QoS requirements associated with real-time applications? For example, User 1's QoS may still benefit from partial completions while User 2 only benefits if all processing is completed. This general class of problems involving both heterogeneous resources and user QoS requirements is the focus of this paper.

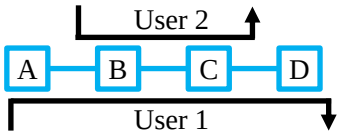

Fig. 1. An example for a network of resources. A, B, C and D represent compute/communication resources. Tasks from User 1 need to be processed on A, B, C, D while tasks from User 2 require processing on B, C.

The design space of possible solutions to this problem is huge and has been explored in many research communities. In this paper we study an approach to resource allocation based on a decomposition of concerns:

(1) user priorities are dynamically set based on the history outcomes;
(2) and, resources are allocated so as to favor users with higher priority.

In such a framework there is quite a bit of latitude in choosing how priorities are set, and in turn how these affect the allocation of resources. For example, users' priorities could be set based on measured deficits, the "difference" of the required and achieved QoS, i.e., Largest Deficit First (LDF) prioritization. In turn, for complex systems such as that in Figure 1, resources could be allocated greedily giving preemptive access to tasks associated with higher-priority users.

In general an optimal user prioritization strategy could leverage detailed information regarding how these priorities will impact the allocation of resources and completion outcomes to achieve the best possible user QoS. Such strategies require excessive amounts of information regarding the underlying compute/communication resources and resource allocation mechanism, and thus are generally hard to implement. By contrast, LDF-based prioritization is quite intuitive. It requires only tracking of users' possibly heterogeneous QoS deficits, in this sense it is truly decoupling user prioritization from the underlying priority-based resource allocation. Unfortunately, it is known to be suboptimal in certain settings [Dimakis and Walrand 2006; Joo et al. 2007; Kang et al. 2013].

A theoretical study of the efficiency and, possibly optimality, of LDF-based prioritization systems supporting real-time users with heterogeneous QoS requirements is the main focus of this paper. We note, however, that we do not directly address the design of the underlying priority-based resource allocation, although we consider some natural characteristics it could have to ensure optimality when combined with LDF user prioritization.

Related Work. There have been much work studying dynamic prioritization policies in the context of diverse resource, workload and/or QoS models.

The authors in [Hou and Kumar 2012; Hou and Kumar 2013] propose a framework to model a wireless access point serving a set of clients that in each period generate packets which need to be transmitted by the end of the period. In their model only one client can transmit at a time and thus the access point can be viewed as a single

resource. Each client transmits its packets over an unreliable channel which has a fixed probability of success, and thus, the time to successfully transmit a packet can be modeled as a geometric random variable. In this setting the authors show that the LDF policy is "optimal." However, the results are restricted to a single resource shared by users with geometric workloads. In this paper we study the performance of LDF in a more general setting which includes this prior work as a special case. This initial set of papers motivated follow-up work in wireless context, see e.g., [Hou and Kumar 2014; Munir et al. 2010; Jaramillo and Srikant 2011].

The performance of LDF and similar policies has also been studied in [McKeown 1995; Dimakis and Walrand 2006; Joo et al. 2007; Kang et al. 2013]. The authors in [Dimakis and Walrand 2006] consider the generalized switch model and were the first to propose the notion of "local pooling" as a sufficient condition for the Longest Queue First (LQF) policy to be throughput optimal. Subsequently, the work in [Joo et al. 2007] considers a multi-hop wireless network under a node-exclusive interference model and shows that the efficiency ratio of the greedy maximal matching policy, which is essentially LQF, equals to the "local pooling" factor of the network graph. More recently, the authors in [Kang et al. 2013] consider real-time traffic in ad hoc wireless networks under a link-interference model and also characterize the efficiency ratio of the LDF policy.

The results in [Dimakis and Walrand 2006; Joo et al. 2007; Kang et al. 2013] depend on the constant service rate model and the specific interference model, i.e., where the set of links/queues that can be scheduled simultaneously is restricted. These models may be appropriate in some wireless/queueing networks but do not necessarily hold in our broader context, e.g., soft real-time applications with stochastic workloads. Also, [Dimakis and Walrand 2006] lacks a performance analysis of LQF when it is not optimal and the works in [Joo et al. 2007; Kang et al. 2013] focus on the efficiency ratio of LDF-like policies but lack a characterization of the full capacity region of these policies. Moreover, when the system can deliver more than the requirements, either the QoS requirements for real-time traffic or the throughput requirements for queueing systems, there is no discussion of how to manage the allocation of the "excess capacity" across users.

The authors in [Tassiulas and Ephremides 1992; Tassiulas and Ephremides 1993; McKeown et al. 1999; Dai and Prabhakar 2000; Stolyar 2004] propose max weight scheduling policies for different types of queueing systems and show them to be throughput optimal via the approaches summarized in [Down and Meyn 1994; Down and Meyn 1997; Meyn and Tweedie 2008]. The authors in [Neely 2009] and [van de Ven et al. 2013] further characterize the delay of the max weight policy, and study its inefficiency in spatial wireless networks, respectively. As we will see in the sequel we too discuss a max weight-like scheduling policy, but it suffers from the usual complexity problems when the decision space is large and it requires excessive amounts of information, motivating us to consider simpler policies.

Additional related work includes work on modeling and scheduling of real-time tasks, see e.g., [Sha et al. 2004; Davis and Burns 2011; Li and Ierapetritou 2008; Shakkottai and Srikant 2002].

Our Contributions. In this paper, we contribute to the theoretical understanding and performance characterization of the Largest Deficit First (LDF) policy with applications to resource allocation to support real-time services. We make three key contributions.

First, we propose a novel general model for a class of systems supporting priority-based resource allocation and study different dynamic prioritization policies. This model is general in terms of the "impact" the priority decisions can have on the QoS payoffs. Specifically, in each period the payoffs under a priority decision are modeled

by a random vector, which includes as special cases the single resource model, the geometric/constant workload and/or specific interference model adopted in prior work. For this general model, we propose a general inner bound $R_{\rm IB}$ for the QoS feasibility region of LDF prioritization policy.

Second, with an additional property, *monotonicity in payoffs*, we characterize the geometry of the inner bound $R_{\rm IB}$. Based on this, we further propose a sufficient condition for the optimality of the LDF policy and characterize the efficiency ratio of LDF. In practice, understanding the geometry of $R_{\rm IB}$ enables us to understand and identify possible bottlenecks in the priority-based resource allocation infrastructure. We also show that the LDF policies (as well as a hierarchical-LDF version) are optimal when there are two classes of exchangeable users.

Finally, we also consider the class of weighted LDF policies, which enable us to explore the allocation of "excess payoffs" when the system has "excess" capacity. Simulation results are exhibited to show the impact of weights and to characterize the clustering of failures.

Paper Organization. The paper is organized as follows: Section 2 introduces our general model for systems supporting priority-based resource allocation. Section 3 develops theoretical results and characterizes the performance of the weighted LDF policies while Section 4 presents some examples for the optimality of the weighted LDF/hierarchical-LDF policies. Section 5 discusses some practical issues while the impact of weights is exhibited via simulation in Section 6. Section 7 concludes the paper and points to future work. Some of the proofs are provided in the Appendix.

2. SYSTEM MODEL

We consider applications which periodically generate random workloads with the same period and specify long-term QoS requirements. In the sequel we let a user denote a specific instance of such an application.

We begin by introducing a general model for systems that allocate resources in each period based on the following decomposition: (1) users are assigned priorities dynamically, e.g., at runtime, according to a function of the past history, and (2) the system allocates resources based on these priorities.

For the most part in this paper, the manner in which (2) is carried out will not be our concern. Instead our focus will be on how to perform dynamic user prioritization to achieve optimal (or near-optimal) system performance when combined with a given underlying mechanism for (2). In our follow-up work [Du and de Veciana 2016], we consider a specific system model and study the combined design of (1) and (2).

2.1. General Model for Systems Supporting Priority-Based Resource Allocation

We consider an abstract system that serves n users indexed from 1 to n. Let $N = \{1, 2, \cdots, n\}$ be the user set. The system operates in discrete time, over periods $t = 1, 2, \cdots$. In each period, it picks a user *priority decision* $\mathbf{d} = (d_1, d_2, \ldots, d_n)$ where d_m is the index of the user with $m^{\rm th}$ highest priority. We let D denote the set of all possible priority decisions and let $|D|$ represent the number of possible decisions, thus, $|D| = n!$

In each period, given the priority decision $\mathbf{d}$ passed to the underlying resources, since there are intrinsic uncertainties in users' workloads, each user i achieves a non-negative random QoS payoff, denoted by $V_i(\mathbf{d})$. We let $\mathbf{V}(\mathbf{d}) = (V_1(\mathbf{d}), V_2(\mathbf{d}), \cdots, V_n(\mathbf{d}))$. We assume the payoffs are independent across periods. The distribution of $\mathbf{V}(\mathbf{d})$ depends on the selected priority decision $\mathbf{d}$ and the expected payoff vector given $\mathbf{d} \in D$ is denoted by $\mathbf{p}(\mathbf{d}) = \mathrm{E}[\mathbf{V}(\mathbf{d})]$. We assume all possible payoff vectors form a finite rational set. Moreover, we naturally assume that for each user $i \in N$, there exists a decision $\mathbf{d}$ such that $p_i(\mathbf{d}) > 0$.

Each user requires a long-term average QoS payoff $q_i \geq 0$ as the QoS requirement. We let $\mathbf{q} = (q_1, q_2, \cdots, q_n)$ and assume q_i's are rational[1]. We denote by $\mathbf{d}(t)$ the priority decision at period t. To keep track of the deficits between required and achieved QoS payoffs, for each user $i \in N$ and period $t+1$, we define[2]

$$X_i(t+1) = [X_i(t) + q_i - V_i(\mathbf{d}(t+1))]^+, \tag{1}$$

where $[x]^+ = \max[x, 0]$.

The goal is thus to devise user prioritization policies which will meet users' long-term payoff requirements.

Definition 2.1. A **user prioritization policy** is a stationary policy that picks a priority decision $\mathbf{d}(t+1) \in D$ at period $t+1$ based on the following:

- users' payoff requirement vector $\mathbf{q}$;
- expected payoff vectors $P = \{\mathbf{p}(\mathbf{d}) | \mathbf{d} \in D\}$;
- and, the deficits $\mathbf{X}(t) = (X_1(t), X_2(t), \cdots, X_n(t))$.

The process $\{\mathbf{X}(t)\}_{t \geq 1}$ is a Markov chain under any such policy. We assume the initial state $\mathbf{X}(0)$, the requirements $\mathbf{q}$, the set of all possible payoff vectors and the user prioritization policy make $\{\mathbf{X}(t)\}_{t \geq 1}$ an irreducible Markov chain.

Definition 2.2. A payoff requirement vector $\mathbf{q}$ is said to be **feasible** if there exists a user prioritization policy η under which the Markov chain $\{\mathbf{X}(t)\}_{t \geq 1}$ is positive recurrent. We also say this policy fulfills this requirement vector.

The expected payoff vectors $P = \{\mathbf{p}(\mathbf{d}) | \mathbf{d} \in D\}$ could in principle be statistically inferred from the history events or by repeated experiments. However, in a practical setting this can be challenging and it is of interest to find a policy that performs well and uses little a-priori information regarding the exponential set of expected payoff vectors P.

Note that this model is general in the sense that the "impact" of priority decisions $\mathbf{d} \in D$ on the QoS payoff vectors P is at this point general, whereas the specific resource and workload models in prior work, e.g., [Dimakis and Walrand 2006; Joo et al. 2007; Kang et al. 2013], implicitly impose properties on P and therefore restrict the results significantly.

2.2. Example: Centralized Computing System for Real-Time Applications

Our model can for example capture a centralized computing infrastructure supporting Soft Real-Time (SRT) applications where the n users share compute resources. In a cloud-based collaborative video conferencing context, a user might correspond to an individual end user and the period length might correspond to the length of a group of video frames.

The users generate streams of tasks periodically. Specifically in each period a user generates several tasks. A task may further consist of a graph of possibly dependent sub-tasks with (possibly) random processing requirements, i.e., workloads. These tasks/sub-tasks need to be fully completed before the end of the period. For real-time services, it is generally useless to process a task after its deadline. For example, in the video conferencing context it is not desirable to present an out-of-date frame. Therefore, we assume tasks/sub-tasks not completed on time are dropped.

[1] All the results in this paper can be generalized to models with irrational values. For simplicity in the proof we do not consider that level of generality.

[2] We truncate the deficit at 0 for the convenience of defining feasibility in the sequel. Removing the truncation won't change the results in the paper.

In each period t, the user prioritization policy picks a user priority decision $\mathbf{d}(t)$, based on which compute resources are allocated to process tasks. Given the task processing results, a payoff $V_i(\mathbf{d}(t))$ is achieved for each user i based on whether the tasks were successfully processed, or how much of the task graphs were completed. In general, $V_i(\mathbf{d}(t))$ may represent any user-specific QoS payoff per period, that can be averaged over time, e.g., the quality/resolution of video frame processing, or the number of task completions. Accordingly the vector $\mathbf{q}$ represents the long-term average QoS requirements.

2.3. Example: Complex Networks and Flexible Modeling of Application Execution Payoffs

As indicated in the introduction, our model also applies to a complex network of heterogeneous compute and communication resources, as long as users periodically and synchronously generate tasks that require timely processing on diverse resources and moving around in the network, e.g., as shown in Figure 1.

Given the priority decision in each period, the network of resources coordinate according to some priority-based resource allocation mechanism to accelerate the processing of tasks with high priorities, by reducing the communication/queueing delays, processing with higher processor speed, allocating more shared resources, etc.

Again, different users can define their payoffs in different ways and specify their QoS requirements accordingly.

3. PERFORMANCE ANALYSIS

In this section we shall develop theoretical results for such systems. Some of these results are similar to prior work but in the more general model while other results are completely new. For completeness we shall develop a self-contained theoretical framework.

3.1. System Feasibility Region and Feasibility Optimal Policy

The set of all feasible long-term payoff requirement vectors will be referred to as the *system feasibility region* F. We let F_η denote the feasibility region of a user prioritization policy η. To characterize F we introduce some further notation.

A vector $\mathbf{x}$ is said to be *dominated* by a vector $\mathbf{y}$ if $x_i \leq y_i$ for all i and is denoted by $\mathbf{x} \preceq \mathbf{y}$. We define $\mathbf{x} \prec \mathbf{y}$, $\mathbf{x} \succeq \mathbf{y}$ and $\mathbf{x} \succ \mathbf{y}$ in a similar manner.

Given the set of priority decisions D and the expected payoff vectors $P = \{\mathbf{p}(\mathbf{d}) | \mathbf{d} \in D\}$, we let C be the set of requirement vectors $\mathbf{q} \in \mathbb{R}_+^n$ which are dominated by a vector in the convex hull of P denoted Conv(P), i.e.,

$$C \equiv \{\mathbf{q} \in \mathbb{R}_+^n \mid \exists \mathbf{x} \in \text{Conv}(P) \text{ such that } \mathbf{q} \preceq \mathbf{x}\}. \tag{2}$$

Figure 2 exhibits C for a two-user (left figure) and three-user (right figure) setting. In the two-user setting, the points labeled $\mathbf{p}(\mathbf{d}_1)$ and $\mathbf{p}(\mathbf{d}_2)$ are the expected payoff vectors of two priority decisions, i.e., where User 1 or User 2 has higher priority, respectively. The shadowed area represents C. In the three-user setting, the circles represent the 6 possible expected payoff vectors, and the region dominated by their convex hull is C. Note that in a n-user scenario where $n \geq 3$, as displayed the expected payoff vectors need not be on a hyperplane in the n-dimensional space. As we will see this is essentially the source of complexity in studying such systems.

Clearly, for any requirement vector $\mathbf{q}$ in the interior of C, denoted by int(C), one can achieve $\mathbf{q}$ if one is allowed to do probabilistic time sharing among priority decisions by picking decisions according to a pre-computed probability distribution whose mean payoff dominates $\mathbf{q}$. Therefore, int(C) $\subseteq F$. We can also show the following result.

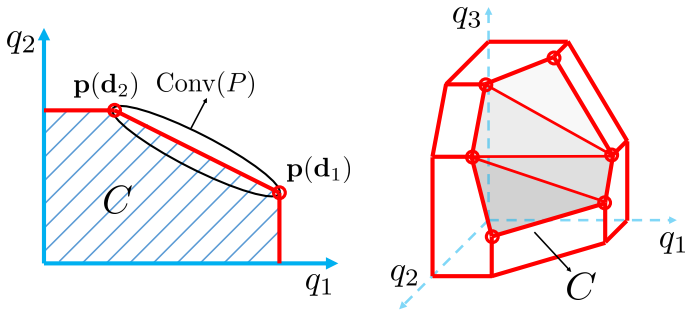

Fig. 2. Examples of set C when $n = 2$ and $n = 3$.

LEMMA 3.1. *The system feasibility region F is such that*

$$F \subseteq \mathbf{cl}(C),$$

where $\mathbf{cl}(C)$ is the closure of C.

Intuitively, if $\mathbf{q}$ is feasible, it is fulfilled by some user prioritization policy that in the long-term picks each priority decision some fraction of the time and thus, $\mathbf{q}$ is dominated by some point in the convex hull of P. This is similar to prior work, e.g., [Tassiulas and Ephremides 1992]. See the appendix for a detailed proof. In other words, C is different from F by at most a boundary, and therefore, characterizes F for practical purposes. Thus, in the sequel we will also refer to C as the system feasibility region.

Ideally, it is desirable to devise an "optimal" policy that can fulfill all feasible requirements. More formally, a user prioritization policy η is said to be *feasibility optimal* if $\mathbf{int}(C) \subseteq F_\eta \subseteq \mathbf{cl}(C)$. Similar to prior work [Tassiulas and Ephremides 1992; Tassiulas and Ephremides 1993], the following max weight-like policy is one such feasibility optimal policy.

Definition 3.2. The **deficit-based max weight (MW)** prioritization policy is such that, at period $t+1$, given the deficit vector $\mathbf{X}(t)$ computed by (1), it picks a priority decision $\mathbf{d}(t+1)$ that satisfies

$$\mathbf{d}(t+1) \in \arg\max_{\mathbf{d} \in D} \langle \mathbf{X}(t), \mathbf{p}(\mathbf{d}) \rangle, \tag{3}$$

where $\langle \mathbf{x}, \mathbf{y} \rangle$ is the inner product of two vectors.

THEOREM 3.3. *The system feasibility region F and the feasibility region of the MW policy F_{MW} are related to C as follows,*

$$\mathbf{int}(C) \subseteq F_{\mathrm{MW}} \subseteq F \subseteq \mathbf{cl}(C),$$

and therefore, the MW policy is feasibility optimal.

See Appendix 8.2 for the proof.

However, the MW policy and time sharing policies require full knowledge of P which is challenging in complex practical systems. Moreover, these policies are hard to implement since they involve solving fairly complex optimization problems, i.e., Eq (3). Changes in the user set or payoff requirement vector $\mathbf{q}$ will also impact the realization of these policies. In summary, the requirements in terms of a-priori knowledge, the computational complexity and lack of flexibility to changes make them hard to use in practice. This motivates the policies considered in the next subsection.

For ease of reference, Table I provides a summary of the notation used to denote various regions used in the rest of the paper—some of these are introduced in the sequel.

Table I. Notation of regions.

Regions	Description
F	System Feasibility Region.
$\text{Conv}(P)$	Convex hull of the expected payoff vectors.
C	Region dominated by $\text{Conv}(P)$.
$F_{\mathbf{w}\text{-LDF}}$	Feasibility region of the w-LDF policy.
R_{IB}	An inner bound for $F_{\mathbf{w}\text{-LDF}}$
B	Dominant of the convex hull.
R	Region characterizing the geometry of R_{IB}.

3.2. Weighted LDF Policies and Associated Feasibility Regions

The LDF user prioritization policies require no a-priori knowledge of the system, are simple to implement and adapt easily to changes in $\mathbf{q}$ or the user set. In particular we shall characterize the feasibility regions of these policies by providing an inner bound.

Definition 3.4. Given a vector $\mathbf{w} = (w_1, w_2, \cdots, w_n) \succ \mathbf{0}$, the **weighted Largest Deficit First (w-LDF)** user prioritization policy is such that, at period $t+1$, given the deficit vector $\mathbf{X}(t)$, it picks a priority decision $\mathbf{d}$ that satisfies

$$w_{d_1} X_{d_1}(t) \geq w_{d_2} X_{d_2}(t) \geq \cdots \geq w_{d_n} X_{d_n}(t),$$

with ties broken arbitrarily (possibly randomly). In other words, it sorts the weighted deficits of users and assigns priorities accordingly. Let $\mathbf{1} \equiv (1, 1, \cdots, 1)$. We refer to the policy with $\mathbf{w} = \mathbf{1}$ the **Largest Deficit First (LDF)** policy.

Clearly, the w-LDF prioritization policies do not require knowledge of the expected payoff vectors P. Note that we still use deficit feedback to stabilize the system. In terms of computational complexity, solving (3) is $O(n!)$ while sorting weighted deficits only requires $O(n \log n)$. It also allows us to further differentiate the performance across users by assigning different weights. The impact of weights is discussed in Section 6.

Prior work has established that the LDF policy need not be feasibility optimal. Therefore, a key question is whether the feasibility regions for the w-LDF policies are acceptable and to characterize the gap between their feasibility regions and the system feasibility region F. To that end, we first provide an inner bound, denoted by R_{IB}, for the feasibility region of any w-LDF policy.

THEOREM 3.5. *For any $\mathbf{w} \succ \mathbf{0}$, an inner bound for the feasibility region of the w-LDF policy $F_{\mathbf{w}\text{-LDF}}$ is given by* $\text{int}(R_{\text{IB}}) \subseteq F_{\mathbf{w}\text{-LDF}}$, *where*

$$R_{\text{IB}} \equiv \{\mathbf{q} \in \mathbb{R}_+^n \mid \exists \boldsymbol{\alpha} \succ \mathbf{0} \text{ such that } \forall S \subseteq N, \sum_{i \in S} \alpha_i q_i \leq \min_{\mathbf{d} \in D(S)} \sum_{i \in S} \alpha_i p_i(\mathbf{d})\} \tag{4}$$

where $D(S)$ denotes the set of all priority decisions that assign the highest $|S|$ priorities to users in S.

In other words, if $\mathbf{q} \in R_{\text{IB}}$, it is feasible under all w-LDF policies except perhaps boundary points. The underlying intuition for this bound is as follows. A vector $\mathbf{q}$ is in R_{IB} if there is a weight vector $\boldsymbol{\alpha} \succ \mathbf{0}$ such that for any subset of users S, and decisions giving users in S the highest priorities, the weighted sum of payoff requirement $\sum_{i \in S} \alpha_i q_i$ will not exceed the least sum weighted payoff $\sum_{i \in S} \alpha_i p_i(\mathbf{d})$. Based on $\boldsymbol{\alpha}$, we can construct

an appropriate Lyapunov function to show feasibility for $\mathbf{q}$ and each $\mathbf{w}$. See Appendix 8.3 for the proof.

Understanding the geometry of R_{IB} enables us to characterize the performance gap between $\mathbf{w}$-LDF and feasibility optimal policies. Let us informally consider the geometry of R_{IB} for the two special cases in Figure 2. In the two-user case in Figure 2, R_{IB} is the same as C and thus, the $\mathbf{w}$-LDF policies are feasibility optimal. However, in the three-user case in Figure 2, this need not be true. Indeed, in this setting, the region R_{IB} corresponds to C minus the convex hull of P, modulo some boundary points. This is exhibited in Figure 3. In the next subsection, we will formalize these observations and show under what conditions they hold true.

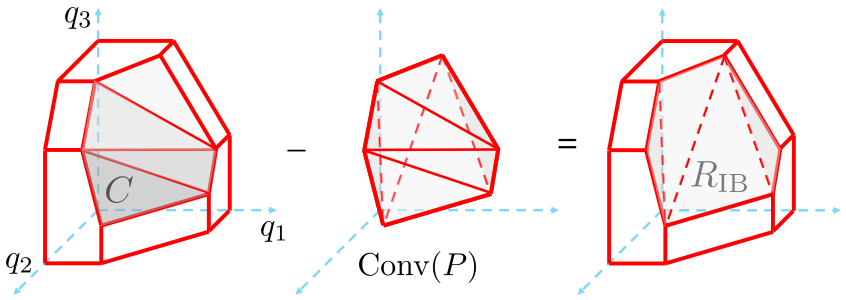

Fig. 3. Visualizing R_{IB} for the three-user scenario in Figure 2. In this example, the expected payoff vectors $P = \{\mathbf{p}(\mathbf{d}) | \mathbf{d} \in D\}$ are not on the same hyperplane and $R_{\text{IB}} = \text{cl}(C - \text{Conv}(P))$.

3.3. Geometry of R_{IB} under Monotonicity in Payoffs

In order to formally characterize the geometry of R_{IB} we will add a further natural requirement to the general model.

We define $S_i(\mathbf{d})$ to be the set of users that have higher priorities than user i under decision $\mathbf{d}$.

Definition 3.6. The system with expected payoff vectors $P = \{\mathbf{p}(\mathbf{d}) | \mathbf{d} \in D\}$ is said to satisfy **monotonicity in individual expected payoff** if, for any two priority decisions $\mathbf{d}_1$ and $\mathbf{d}_2$ and any user i such that $S_i(\mathbf{d}_1) \subseteq S_i(\mathbf{d}_2)$, it is true that $p_i(\mathbf{d}_1) \geq p_i(\mathbf{d}_2)$. We call this **monotonicity in payoffs** for short.

In other words, a user i can expect to get a higher payoff if some users with higher priority are re-assigned lower priorities. This property characterizes in a broad sense how priorities impact the expected payoffs when the underlying system allocates resources. It is a natural condition but need not hold in general.

We shall define B to be the set of payoff requirement vectors $\mathbf{q}$ which dominate a vector in the convex hull of P, i.e.,

$$B \equiv \{\mathbf{q} \in \mathbb{R}_+^n \mid \exists \mathbf{x} \in \text{Conv}(P) \text{ such that } \mathbf{q} \succeq \mathbf{x}\}.$$

We call B the *dominant of the convex hull*. Contrast this to the definition of C in (2).

For the special cases in Figure 2 and 3, $B \cap C$ equals to $\text{Conv}(P)$, but in general it can be larger than $\text{Conv}(P)$. Figure 4 shows a conceptual picture of what could happen. The three circles represent three possible expected payoff vectors. Here, the whole shadowed area $B \cap C$ is larger than the region $\text{Conv}(P)$ which is the triangle formed by the three circles. Note that this is only a conceptual example to help visualize $B \cap C$ in higher dimensions. In reality for two dimensions, i.e., systems with two users, we know there are only 2 expected payoff vectors as shown in Figure 2.

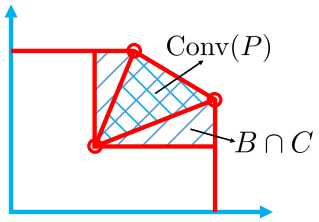

Fig. 4. An example where $B \cap C$ is larger than $\text{Conv}(P)$.

In the sequel we will see that given monotonicity in payoffs, R_{IB} is obtained by "removing" $B \cap C$, rather than just $\text{Conv}(P)$ from C. To develop this result we need some further notation associated with each subset of users $S \subseteq \{1, 2, \cdots, n\}$.

The projection of a vector $\mathbf{x}$ on the subspace of S is denoted by $\mathbf{x}^S$, i.e.,

$$x_i^S = \begin{cases} x_i & \text{if } i \in S \\ 0 & \text{otherwise.} \end{cases}$$

We let $P^S \equiv \{\mathbf{p}^S(\mathbf{d}) | \mathbf{d} \in D(S)\}$ represent the projections of expected payoff vectors corresponding to decisions in $D(S)$, i.e., which assign the highest priorities to users in S.

Given a subset S and P^S, we define the feasibility region C^S and the dominant of the convex hull B^S as follows.

$$C^S \equiv \{\mathbf{q}^S \in \mathbb{R}_+^n \mid \exists \mathbf{x}^S \in \text{Conv}(P^S) \text{ such that } \mathbf{q}^S \preceq \mathbf{x}^S\},$$
$$B^S \equiv \{\mathbf{q}^S \in \mathbb{R}_+^n \mid \exists \mathbf{x}^S \in \text{Conv}(P^S) \text{ such that } \mathbf{q}^S \succeq \mathbf{x}^S\}.$$

Note that C^S and B^S are not necessarily the same as projecting C and B on the subspace of S, respectively. This is because in the definitions of C^S and B^S, we only focus on a subset of decisions $D(S)$ rather that the full decision set D.

Let us now define a region R which will help characterize the geometry of the inner bound R_{IB}.

Definition 3.7. Let R be defined as follows:

$$R \equiv \{\mathbf{q} \in \mathbb{R}_+^n \mid \forall S \subseteq N, \mathbf{q}^S \in C^S \setminus B^S\},$$

where $C^S \setminus B^S = \{\mathbf{q}^S | \mathbf{q}^S \in C^S, \mathbf{q} \notin B^S\}$. In other words, any $\mathbf{q} \in R$ satisfies that for any user subset S, its projection on the subspace of S belongs to the set $C^S \setminus B^S$, which is the feasibility region C^S minus the dominant of the convex hull B^S.

One can visualize obtaining the set R as a process of removing $B^S \cap C^S$ from C^S in all subspaces corresponding to all subsets S. The geometry of R_{IB} is then captured as follows.

THEOREM 3.8. *If the system satisfies monotonicity in payoffs, then the inner bound region R_{IB} is such that*

$$\text{int}(R) \subseteq R_{\text{IB}} \subseteq \text{cl}(R).$$

See Appendix 8.4 for this somewhat intricate argument.

3.4. Sufficient Condition for w-LDF's Optimality

By Theorem 3.5 and Theorem 3.8, we immediately get

$$\text{int}(R) \subseteq \text{int}(R_{\text{IB}}) \subseteq F_{\text{w-LDF}}. \tag{5}$$

Since R is obtained by removing $B^S \cap C^S$ from C^S for each S, if what is removed is nothing more than a boundary, the difference between R and C is at most a boundary and thus w-LDF policies are feasibility optimal. It is easy to see this happens when vectors in P^S lie on a hyperplane for each subset of users S. This can be formalized as follows.

Definition 3.9. The system with expected payoff vectors $P = \{\mathbf{p}(\mathbf{d})|\mathbf{d} \in D\}$ is said to satisfy **subset payoff equivalence** if for each subset of users S the vectors in $P^S = \{\mathbf{p}^S(\mathbf{d})|\mathbf{d} \in D(S)\}$ lie on a hyperplane, i.e., there exists a nonzero $\boldsymbol{\alpha}^S \succeq \mathbf{0}$ such that for all $\mathbf{d}_1, \mathbf{d}_2 \in D(S)$,

$$\langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}_1)\rangle = \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}_2)\rangle.$$

THEOREM 3.10. *If the system satisfies monotonicity in payoffs and subset payoff equivalence, then*

$$\mathrm{int}(C) \subseteq F_{\mathbf{w}\text{-LDF}} \subseteq \mathbf{cl}(C),$$

and therefore, the w*-LDF policies are feasibility optimal.*

Please refer to Appendix 8.7 for detailed proof.

The conditions for this theorem are akin but not equivalent to the conditions introduced in [Dimakis and Walrand 2006] for the generalized switch model. Specifically, we require the system to satisfy monotonicity in payoffs and subset payoff equivalence. The work in [Dimakis and Walrand 2006] requires local pooling in the generalized switch model. In the model in [Dimakis and Walrand 2006], given a priority decision $\mathbf{d} = (d_1, d_2, \cdots, d_n)$ where d_i is the index of the queue with the the i^{th} highest priority, the queue service rate vector can be denoted by $\mathbf{m}(\mathbf{d})$, where $m_{d_i}(\mathbf{d})$ represents the units of work that can be removed from queue d_i in one time slot under priority decision $\mathbf{d}$. $\mathbf{m}(\mathbf{d})$ is akin to $\mathbf{p}(\mathbf{d})$ in our context. However, the generalized switch model in [Dimakis and Walrand 2006] implies properties on the service rate vectors $\mathbf{m}(\mathbf{d})$. For example, it implies that for all $\mathbf{d} = (d_1, d_2, \cdots, d_n)$, we have $m_{d_1}(\mathbf{d}) \geq m_{d_1}(\mathbf{d}')$ for all $\mathbf{d}'$, and $m_{d_2}(\mathbf{d}) \geq m_{d_2}(\mathbf{d}')$ for all $\mathbf{d}'$ satisfying $m_{d_1}(\mathbf{d}) = m_{d_1}(\mathbf{d}')$, etc. These implicit requirements do not necessarily hold in systems which satisfy the conditions in Theorem 3.10.

If the system has only two users, then clearly subset payoff equivalence is satisfied since the two expected payoff vectors are always on a line. Therefore, we get the following corollary.

COROLLARY 3.11. *If the system has two users and satisfies monotonicity in payoffs, then* w*-LDF policies are feasibility optimal.*

Note that in a two-user scenario, the property of monotonicity in payoffs simply means a user gets higher payoff under the higher priority than its payoff under the lower priority. In Section 4.2 we will consider systems serving two classes of exchangeable users and use this corollary to show the optimality of LDF-like policies.

3.5. Efficiency Ratio Analysis

When the conditions in Theorem 3.10 do not hold, one can still study the efficiency ratio, see e.g., [Joo et al. 2007], to evaluate the performance of w-LDF policies.

Definition 3.12. The **efficiency ratio** of the w-LDF policy is defined as

$$\gamma_{\mathbf{w}\text{-LDF}} = \sup\{\gamma|\gamma F \subseteq F_{\mathbf{w}\text{-LDF}}\}.$$

Clearly $\gamma_{\mathbf{w}\text{-LDF}}$ equals to 1 if and only if the w-LDF policy is feasibility optimal.

If a system does not satisfy subset payoff equivalence, i.e., for some subset of users S the vectors in P^S are not on the same hyperplane, we can characterize the "heterogeneity" of these vectors based on the following notion.

Definition 3.13. Given a subset of users $S \subseteq N$, the **subset payoff ratio** σ_S for S is defined as

$$\sigma_S = \max_{\substack{\boldsymbol{\alpha}^S \succeq \mathbf{0} \\ \boldsymbol{\alpha}^S \neq \mathbf{0}}} \frac{\min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle}{\max_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle}. \tag{6}$$

The optimal $\boldsymbol{\alpha}^S$ is such that the projections of the vectors in P^S on $\boldsymbol{\alpha}^S$ are as close to each other as possible.

Clearly if the vectors in P^S are on the same hyperplane, then $\sigma_S = 1$ and the optimal $\boldsymbol{\alpha}^S$ is the normal vector to the hyperplane. Intuitively, σ_S characterizes the degree to which the vectors in P^S deviate from being on the same hyperplane.

This notion enables us to characterize the efficiency ratio of w-LDF for a given system.

THEOREM 3.14. *If the system satisfies monotonicity in payoffs, the efficiency ratio of the* w*-LDF policy is such that*

$$\gamma_{\textbf{w-LDF}} \geq \min_{S \subseteq N} \sigma_S.$$

See Appendix 8.8 for the proof. Intuitively, the bottleneck of the efficiency ratio is the subset S where σ_S is the smallest.

Note that by picking any $\boldsymbol{\alpha} \succ \mathbf{0}$, we can get lower bounds on σ_S for all subsets $S \subseteq N$ by placing its projection $\boldsymbol{\alpha}^S$ into (6). Thus, any $\boldsymbol{\alpha} \succ \mathbf{0}$ enables us to construct a lower bound on $\gamma_{\text{w-LDF}}$. A trivial option is $\boldsymbol{\alpha} = \mathbf{1}$, where for each subset S the value of $\langle \mathbf{1}^S, \mathbf{p}^S(\mathbf{d}) \rangle$ represents the sum payoff of users in S under decision $\mathbf{d}$.

We have shown that the efficiency and optimality of the w-LDF policies is related to R_{IB}. Understanding and analyzing the geometry of R_{IB} can in principle enable us to provide feedback to the designers of priority-based resource allocation mechanisms regarding which specific priority decision or set of priority decisions are problematic and bottlenecks for the system so that the designers can focus on improving the resource allocation for these problematic decisions. For example, in the conceptual setting shown in Figure 4, the priority decision corresponding to the lower left circle is the "bottleneck" of the system and should be targeted to make the dominant of the convex hull as small as possible. This is of particular interest for some practical systems where it is possible to get explicit knowledge of P which reflect the underlying priority-based resource allocation, e.g., by collecting data over a long time.

A priority decision is problematic if the associated underlying resource allocation suffers from resource contention, blocking among users/applications, or even deadlocks on compute resources, etc. Based on feedback regarding the bottlenecks, the designer could improve the associated resource allocation schemes, e.g., by increasing the processing speed of the certain computing resources, spending more energy, reducing the contention, and/or resolving the blocking/deadlock, and thus, improve the efficiency of the overall system under the w-LDF prioritization policies.

4. EXAMPLES FOR W-LDF'S OPTIMALITY

Theorem 3.10 gives a sufficient condition for w-LDF to be feasibility optimal. One example system that satisfies these conditions is the model considered in prior work [Hou and Kumar 2012] which, as mentioned in Section 1, can be viewed as a single-

resource geometric-workload model. In this section we consider more system settings and show how our results provide useful insights in practice.

4.1. Exchangeable Expected Payoffs

We shall start by showing that for systems that are "symmetric", w-LDF policies are feasibility optimal.

Definition 4.1. A subset of users S is said to have **exchangeable expected payoffs** if, for all priority decisions $\mathbf{d} \in D$ and all $i, j \in S$, if we switch the priorities of user i and j and use $\mathbf{d}'$ to represent the resulting new priority decision, then

$$p_k(\mathbf{d}') = \begin{cases} p_k(\mathbf{d}) & \text{if } k \neq i, j \\ p_j(\mathbf{d}) & \text{if } k = i \\ p_i(\mathbf{d}) & \text{if } k = j. \end{cases}$$

In other words, exchanging the priorities of two users in S will simply exchange their expected payoffs without impacting that of other users. This would be true if the priority-based resource allocation were symmetric for users in S and the users generate tasks with identically distributed or exchangeable workloads.

If the users in N have exchangeable expected payoffs, we can verify the property of subset payoff equivalence by picking $\boldsymbol{\alpha}^S = \mathbf{1}^S$ for each subset of users S. Therefore, by Theorem 3.10 we get the following corollary.

COROLLARY 4.2. *If the set of users N have exchangeable expected payoffs and the system satisfies monotonicity in payoffs, then the* w*-LDF policies are feasibility optimal.*

See Appendix 8.9 for the proof.

4.2. Multiple Classes of Exchangeable Users and Hierarchical-LDF

In this subsection, we first consider a system supporting two classes of exchangeable users. Formally, a class of users is *exchangeable* if they have exchangeable expected payoffs and the same QoS requirement. The users in different classes may have distinct payoffs and QoS requirements. In some contexts it is of practical interest to first prioritize the classes and then prioritize users in each class, respectively. We refer to such schemes as using *class-based hierarchical prioritization*.

In practice, depending on whether the priorities of classes can change dynamically, there are two types of class-based hierarchical prioritization: Type 1 where the class priorities are fixed, and Type 2 where one can dynamically prioritize classes of users, and then users within each class.

The first type of hierarchical prioritization might correspond to a setting where the users/applications are separated into human-interactive/high-QoS and background-processing/low-QoS categories [Gatherer 2015], and it is always desirable to first process high-QoS users. In this setting, the problem is reduced to a collection of independent user prioritization problems similar to the one considered in this paper. By Corollary 4.2, w-LDF is feasibility optimal to prioritize users in each class.

The second type of dynamic hierarchical prioritization might be of interest in systems where switching between processing different user classes involves overheads, and/or where it is inefficient to mix the processing of different user classes, probably because of resource contention or deadlocks.

In this setting, we propose a class-based *hierarchical-LDF* policy that in each period works in two steps by (1) prioritizing classes by LDF based on the aggregate deficits, i.e., the sum of deficits for users in the same class, and (2) prioritizing users in each class according to LDF based on individual users' deficits. The framework of

hierarchical-LDF is exhibited in Figure 5. Note that here LDF can be replaced by w-LDF for any $\mathbf{w} \succ \mathbf{0}$ and the following result would hold.

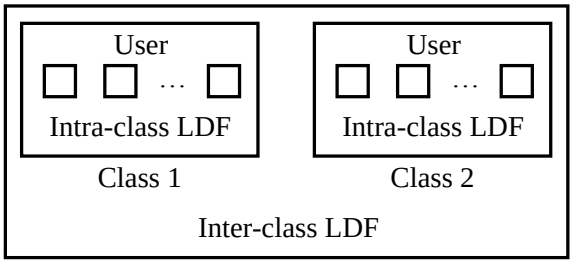

Fig. 5. The framework for class-based hierarchical-LDF policy.

THEOREM 4.3. *In a system with two classes of exchangeable users, if the property of monotonicity in payoffs is satisfied, the hierarchical-LDF policy is feasibility optimal among all possible class-based hierarchical prioritization policies.*

The proof follows directly from Corollary 3.11 and 4.2. By Corollary 3.11 we know the class-based LDF policy is optimal to set priorities amongst the two classes and by Corollary 4.2 we know the LDF-based user prioritization is also optimal for the exchangeable users in each class.

More generally, for systems serving multiple (more than two) classes of exchangeable users, one can view each class as a "super user", and define the aggregate payoff and QoS requirement for a super user to be the sum of payoffs and QoS requirements for users in that class, respectively. Then the dynamic prioritization of super users can be viewed as the problem considered in this paper. Therefore, by Theorem 3.10, if the system with the super users' expected aggregate payoffs satisfies monotonicity in payoffs and subset payoff equivalence, the LDF policy is a feasibility optimal choice for prioritizing super users and thus, the hierarchical-LDF policy is feasibility optimal among all class-based hierarchical prioritization policies. Indeed, all the results we have introduced, e.g., Theorem 3.3-3.14, still hold for the prioritization of these super users.

5. SOME PRACTICAL ISSUES

In practice, besides meeting minimum payoff requirements, users may be willing to pay for additional payoffs, e.g., better video quality in the video conferencing setting, albeit at possibly different prices. Given the requirements $\mathbf{q}$ and the achieved average payoffs $\mathbf{p} = (p_1, p_2, \cdots, p_n)$, we call $p_i - q_i$ the *excess payoff* for each user i. While using w-LDF policies to fulfill users' payoff requirements, we also want to manage the allocation of excess payoffs across users, perhaps with the aim of maximizing the benefits to the system or users.

However, the non-negative definition of deficit (1) makes it hard to track excess payoffs. For example, consider a model with 2 users and suppose the payoff is always 1 for the high priority user and 0 for the low priority user. Suppose the payoff requirement vector is $\mathbf{q} = (0.1, 0.5)$. Since $1 > 0.1 + 0.5$, we know $\mathbf{q}$ is feasible and the system can deliver 0.4 excess payoff. Suppose we use the LDF policy, starting from $\mathbf{X}(0) = (0, 0)$ it is easy to verify[3] that the system will switch giving high priority to these two users, and

[3]Since the payoffs are deterministic, we can verify this by evaluating the deficits for the first few periods and we will observe that the process $\{\mathbf{X}(t)\}_{t\geq 1}$ evolves in a periodic pattern.

thus the achieved average payoff vector is $\mathbf{p} = (0.5, 0.5)$. Clearly User 1 gets 0.4 excess payoff while User 2 gets nothing. This happens because $X_1(t)$ and $X_2(t)$ are frequently forced to 0 from different negative values, which causes the "unfairness" between these two users.

To solve this problem, we modify the deficit definition for each user i and period $t+1$ as follows,

$$X_i'(t+1) = X_i'(t) + q_i - V_i(\mathbf{d}(t+1)), \tag{7}$$

i.e., we allow $X_i'(t)$ to be negative.

Now for the simple example above, if we adopt LDF but based on the possibly negative deficits $\mathbf{X}'(t) = (X_1'(t), X_2'(t), \cdots, X_n'(t))$, we can get achieved average payoff vector $\mathbf{p} = (0.3, 0.7)$. We observe that the two users equally split the excess payoff.

Intuitively, for each user i the modified deficit $X_i'(t)$ changes roughly linearly as t increases with the slope being $q_i - p_i$. Since w-LDF policy aims to balance weighted deficit $w_i X_i'(t)$, we know $w_i(p_i - q_i)$ is roughly the same for all users. We will verify this observation in the simulation section and based on this we can manage the excess payoffs across users by picking the appropriate weight vector $\mathbf{w}$.

Note that for completeness we will need to modify the feasibility definition since the process $\{\mathbf{X}'(t)\}_{t\geq 1}$ is no longer positive recurrent as it may keep decreasing or increasing. Now we call a payoff requirement vector $\mathbf{q}$ feasible if, under some user prioritization policy, for each user i the time-averaged payoff per period is at least q_i. Formally, recall that $V_i(\mathbf{d}(t))$ is the random payoff for user i in period t. As the payoff requirement, each user i requires that

$$\liminf_{\tau\to\infty} \frac{1}{\tau} \sum_{t=1}^{\tau} V_i(\mathbf{d}(t)) \geq q_i, \text{with probability 1}.$$

Note that this definition and Definition 2.2 are just two ways to define the feasibility. With the theorem in [Blackwell 1956] we can show that for any user prioritization policy, the sets of feasible QoS requirements under these two different feasibility definitions differ by at most a boundary and thus are equivalent for practical purposes. Therefore, all the results we discuss in this paper hold under both feasibility definitions.

6. SIMULATIONS

In this section we explore via simulation the impact of weights of w-LDF policies.

Consider an illustrative system with single computing resource serving 3 soft real-time users. In each period of length $\delta = 10$, each user generates one task that need to complete by end of the period. We let the non-negative workload, i.e., task service time, distributions for three users be Gamma$(12, 0.5)$, Gamma$(4, 1)$ and Gamma$(10, 0.1)$, respectively. We pick these workload distributions to make them general and heterogeneous. In each period, the payoff for user i is 1 if user i's task completes and is 0 otherwise. Accordingly, user i's QoS requirement q_i represents the long-term task completion ratio.

We start with initial deficit $\mathbf{X}'(0) = (0, 0, 0)$. In each period, we independently generate task workloads for users and simulate the w-LDF policy based on $\mathbf{X}'(t)$ to pick a priority decision. The single resource sequentially processes users' tasks from highest to lowest priority. Tasks not completed on time are dropped. All simulations are run for 30000 periods. A requirement vector $\mathbf{q}$ is feasible if it is dominated by the achieved task completion ratio vector $\mathbf{p}$ over the 30000 periods. The vectors $\mathbf{q}$ and $\mathbf{w}$ are specified in various settings in the sequel.

Note that in this setting monotonicity in payoffs is satisfied while subset payoff equivalence is not.

6.1. Impact of Weights on Long-Term Completion Ratios

In Table II we consider a requirement vector $\mathbf{q}$ that is feasible under the $\mathbf{w}$-LDF policies and display the achieved $\mathbf{p}$ under two different weight vectors $\mathbf{w}$. For each weight vector $\mathbf{w}$, we verify that $w_i(p_i - q_i)$ is the same for all three users. Contrasting the two lines in Table II, we can see that for a system which can deliver more than required, changing the weight vector reallocates the excess payoffs and gives more excess payoffs to users with smaller weights.

Table II. Achieved completion ratio vectors under two weight vectors.

$\mathbf{q}$	$\mathbf{w}$	Achieved $\mathbf{p}$	$w_i(p_i - q_i)$
$0.8, 0.6, 0.4$	$(1, 1, 1)$	$0.85, 0.65, 0.45$	0.05
	$(10, 1, 1)$	$0.809, 0.69, 0.49$	0.09

6.2. Characterization of Clustering of Failures and Impact of Weights

If a user's task is not completed in a period, we call it a failure event. The requirement vector $\mathbf{q}$ focuses on long-term task completion ratio, but it would likely be undesirable for a user to experience consecutive or clustered failure events. Figure 6 gives an example of failure events. In this subsection we consider the same $\mathbf{q} = (0.8, 0.6, 0.4)$ used above and explore the clustering of failures under two $\mathbf{w}$-LDF policies.

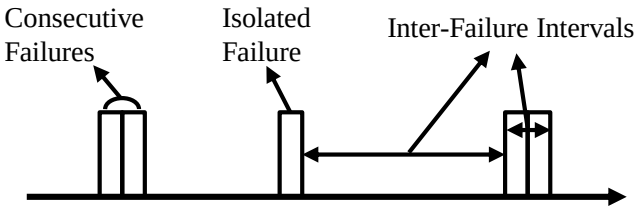

Fig. 6. Characteristics of clustering of failures.

We consider Inter-Failure Intervals (IFIs) between typical failures. IFI is supported on the set $\{1, 2, 3, \cdots\}$. To quantitatively evaluate the clustering of the failures, we focus on the standard deviation (SD) of the IFIs for each user. One extreme case is that failures happen strictly periodically and therefore, the SD is 0. Intuitively, a user with a smaller IFI SD implies that the user experiences less clustered failures.

Next we introduce an evaluation benchmark. For each user i, we know $1 - p_i$ represents the time-averaged failure ratio. If the failure happens in each period independently with probability $1 - p_i$, the IFI can be modeled by a geometric random variable supported on the set $\{1, 2, 3, \cdots\}$ with the parameter being $1 - p_i$. We use the SD of such a geometric random variable as a benchmark.

Under some $\mathbf{w}$-LDF policy, we define *SD ratio* of user i to be the ratio of user i's IFI SD to the SD of the geometric random variable with parameter $1 - p_i$. Table III shows the SD ratios of three users under two different weight vectors $\mathbf{w}$. Under $\mathbf{w} = (1, 1, 1)$, the ratios are less than 1, indicating that the failures under the LDF policy are less clustered compared to the scheme where failure event happens i.i.d. in each period. The last two columns in Table III indicates that increasing the weight of user i reduces the degree of failure clustering for user i but at the price of other users' more clustered

failures. Thus, the users' sensitivities to clustered failures is another factor to consider when one assigns weights to users.

Table III. Characterization of clustering of failures.

	SD ratio under $\mathbf{w} = (1, 1, 1)$	**SD ratio under** $\mathbf{w} = (10, 1, 1)$
User 1	88%	39%
User 2	77%	97%
User 3	92%	107%

7. CONCLUSION

Resource allocation in complex systems supporting real-time users with general QoS requirements can be relatively "easy". One can in principle design the system to allow priority-based resource allocation and adopt simple w-LDF policies to dynamically prioritize users/applications. Our theory provides guidance towards understanding the suboptimality and even optimality of such solutions and how to improve the system design. For future work, it would be interesting to explore the management of real-time users across systems and/or sharing with non real-time traffic.

REFERENCES

David Blackwell. 1956. An Analog of the minimax theorem for vector payoffs. *Pacific J. Math.* 6, 1 (November 1956), 1–8.

Stephen Boyd and Lieven Vandenberghe. 2009. *Convex Optimization*. Cambridge university press.

J.H. Conway and N.J.A. Sloane. 2013. *Sphere Packings, Lattices and Groups*. Springer.

J.G. Dai and Balaji Prabhakar. 2000. The throughput of data switches with and without speedup. In *Proceedings of INFOCOM 2000*. 556–564.

Robert I. Davis and Alan Burns. 2011. A Survey of Hard Real-Time Scheduling for Multiprocessor Systems. *Comput. Surveys* 43 (October 2011). Issue 4.

Antonis Dimakis and Jean Walrand. 2006. Sufficient Conditions for Stability of Longest-Queue-First Scheduling: Second-Order Properties Using Fluid Limits. *Advances in Applied Probability* 38, 2 (June 2006).

D. Down and S. Meyn. 1994. A Survey of Markovian Methods for Stability of Networks. In *11th International Conference on Analysis and Optimization of Systems*.

D. Down and S.P. Meyn. 1997. Piecewise linear test functions for stability and instability of queueing networks. *Queueing Systems* 27 (April 1997), 205–226. Issue 3-4.

Yuhuan Du and Gustavo de Veciana. 2016. Scheduling for Cloud-Based Computing Systems to Support Soft Real-Time Applications. *INFOCOM 2016* (April 2016).

Alan Gatherer. 2015. Personal communication. (February 2015).

I-Hong Hou and P. R. Kumar. 2012. Queueing systems with hard delay constraints: a framework for real-time communication over unreliable wireless channels. *Queueing Systems* 71 (March 2012), 151–177. Issue 1-2.

I-Hong Hou and P. R. Kumar. 2013. *Packets with Deadlines: A Framework for Real-Time Wireless Networks*. Morgan & Claypool Publishers.

I-Hong Hou and P. R. Kumar. 2014. Scheduling Heterogeneous Real-Time Traffic over Fading Wireless Channels. *IEEE/ACM Transactions on Networking* 22 (October 2014), 1631–1644. Issue 5.

Juan Jose Jaramillo and R. Srikant. 2011. Optimal Scheduling for Fair Resource Allocation in Ad Hoc Networks With Elastic and Inelastic Traffic. *IEEE Transactions on Networking* 19 (August 2011), 1125–1136. Issue 4.

Changhee Joo, Xiaojun Lin, and Ness B. Shroff. 2007. Performance Limits of Greedy Maximal Matching in Multi-hop Wireless Networks. In *IEEE Conference on Decision and Control*. 1128–1133.

Xiaohan Kang, Weina Wang, Juan Jose Jaramillo, and Lei Ying. 2013. On the Performance of Largest-Deficit-First for Scheduling Real-Time Traffic in Wireless Networks. In *Proceedings of MobiHoc*. 99–108.

Zukui Li and Marianthi Ierapetritou. 2008. Process scheduling under uncertainty: Review and challenges. *Computers and Chemical Engineering* 32 (2008), 715–727. Issue 4-5.

Nicholas McKeown. 1995. Scheduling Algorithms for Input-Queued Cell Switches. (1995). Ph.D. dissertation.

Nick McKeown, Adisak Mekkittikul, Venkat Anantharam, and Jean Walrand. 1999. Achieving 100% Throughput in an Input-Queued Switch. *IEEE Transactions on Communications* 47, 8 (August 1999), 1260–1267.

S. P. Meyn and R. L. Tweedie. 2008. *Markov Chains and Stochastic Stability*. Cambridge University Press.

Sirajum Munir, Shan Lin, Enamul Hoque, S. M. Shahriar Nirjon, J. A. Stankovic, and K. Whitehouse. 2010. Addressing Burstiness for Reliable Communication and Latency Bound Generation in Wireless Sensor Networks. In *IPSN 2010*. 303–314.

M. J. Neely. 2009. Delay Analysis for Max Weight Opportunistic Scheduling in Wireless Systems. *IEEE Trans. Automat. Control* 54 (September 2009), 2137–2150. Issue 9.

Lui Sha, Tarek Abdelzaher, Karl-Erik ÅRZÉN, Anton Cervin, Theodore Baker, Alan Burns, Giorgio Buttazzo, Marco Caccamo, John Lehoczky, and Aloysius K. Mok. 2004. Real Time Scheduling Theory: A Historical Perspective. *Real-Time Systems* 28 (November-December 2004), 101–155. Issue 2-3.

Sanjay Shakkottai and R. Srikant. 2002. Scheduling Real-Time Traffic With Deadlines over a Wireless Channel. *Wireless Networks* 8 (January 2002), 13–26. Issue 1.

Alexander L. Stolyar. 2004. Maxweight Scheduling in a Generalized Switch: State Space Collapse and Workload Minimization in Heavy Traffic. *The Annals of Applied Probability* 14, 1 (February 2004).

Leandros Tassiulas and Anthony Ephremides. 1992. Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks. *IEEE Trans. Automat. Control* 37 (December 1992), 1936–1948. Issue 12.

Leandros Tassiulas and Anthony Ephremides. 1993. Dynamic Server Allocation to Parallel Queues with Randomly Varying Connectivity. *IEEE Transactions on Information Theory* 39, 2 (March 1993), 466–478.

P.M. van de Ven, S.C. Borst, and L. Ying. 2013. Inefficiency of MaxWeight scheduling in spatial wireless networks. *Computer Communications* 36 (July 2013), 1350–1359. Issue 12.

8. APPENDIX

8.1. Proof of Lemma 3.1

Given $\mathbf{q} \in F$, since it is feasible there exists a user prioritization policy η that fulfills $\mathbf{q}$, i.e., the Markov chain $\{\mathbf{X}(t)\}_{t\geq 1}$ is positive recurrent, which implies there exists a stationary distribution over the state space. Since η is a stationary policy that picks $\mathbf{d}(t+1)$ based on $\mathbf{X}(t)$, by Ergodic Theorem each priority decision $\mathbf{d}$ is selected with some time fraction $\alpha_{\mathbf{d}}$ such that $\sum_{\mathbf{d}\in D} \alpha_{\mathbf{d}} = 1$. If we consider $X_i(t)$ as a queue, the average arrival q_i should not be bigger than the average departure which is given by $\sum_{\mathbf{d}\in D} \alpha_{\mathbf{d}} p_i(\mathbf{d})$ since otherwise $X_i(t)$ goes a.s. to infinity and the chain cannot be positive recurrent.

Therefore,

$$\mathbf{q} \preceq \sum_{\mathbf{d}\in D} \alpha_{\mathbf{d}} \mathbf{p}(\mathbf{d}) \in \mathrm{Conv}(P),$$

which implies $\mathbf{q} \in C \subseteq \mathbf{cl}(C)$.

8.2. Proof of Theorem 3.3

We start by introducing a lemma.

LEMMA 8.1. *A payoff requirement vector $\mathbf{q}$ is in C if and only if, for any non-negative vector $\boldsymbol{\gamma} \succeq 0$, there exists a priority decision $\mathbf{d}$, such that*

$$\langle \boldsymbol{\gamma}, \mathbf{q} \rangle \leq \langle \boldsymbol{\gamma}, \mathbf{p}(\mathbf{d}) \rangle.$$

Equivalently, $\mathbf{q}$ *is in* C *if and only if, for any* $\boldsymbol{\gamma} \succeq \mathbf{0}$*,*

$$\langle \boldsymbol{\gamma}, \mathbf{q} \rangle \leq \max_{\mathbf{d} \in D} \langle \boldsymbol{\gamma}, \mathbf{p}(\mathbf{d}) \rangle.$$

To understand this lemma, since any vector in C is dominated by some $\mathbf{q} \in \text{Conv}(P)$, we consider for simplicity a $\mathbf{q} \in \text{Conv}(P)$. Such a vector can be expressed as a convex combination of expected payoff vectors, i.e., $\mathbf{q} = \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} \mathbf{p}(\mathbf{d})$, where $\sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} = 1$ and $\alpha_{\mathbf{d}} \geq 0, \forall \mathbf{d} \in D$.

Now we have

$$\begin{aligned} \langle \boldsymbol{\gamma}, \mathbf{q} \rangle &= \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} \langle \boldsymbol{\gamma}, \mathbf{p}(\mathbf{d}) \rangle \\ &\leq \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} \max_{\mathbf{d} \in D} \langle \boldsymbol{\gamma}, \mathbf{p}(\mathbf{d}) \rangle \\ &= \max_{\mathbf{d} \in D} \langle \boldsymbol{\gamma}, \mathbf{p}(\mathbf{d}) \rangle. \end{aligned}$$

This actually proves the necessity of the condition. The formal proof is shown below.

PROOF PROOF OF LEMMA 8.1. Given a payoff requirement vector $\mathbf{q}$, by definition of C, we know that $\mathbf{q}$ lying in set C is equivalent to the *feasibility* of the following set of linear equations and inequalities,

$$\begin{cases} \mathbf{q} \preceq \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} \mathbf{p}(\mathbf{d}) \\ \alpha_{\mathbf{d}} \geq 0, \forall \mathbf{d} \in D \\ \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} = 1. \end{cases} \tag{8}$$

The condition in Lemma 8.1 is equivalent to the *infeasibility* of

$$\begin{cases} \boldsymbol{\gamma} \succeq \mathbf{0} \\ \langle \boldsymbol{\gamma}, \mathbf{q} \rangle > \langle \boldsymbol{\gamma}, \mathbf{p}(\mathbf{d}) \rangle, \forall \mathbf{d} \in D. \end{cases} \tag{9}$$

By strong duality [Boyd and Vandenberghe 2009] it is easy to prove that set (8) being feasible is equivalent to set (9) being infeasible and this concludes the proof. □

PROOF PROOF OF THEOREM 3.3. By Lemma 3.1 we know $F \subseteq \text{cl}(C)$ and by definition we know $F_{\text{MW}} \subseteq F$. To prove the theorem it suffices to show $\text{int}(C) \subseteq F_{\text{MW}}$. We show this by constructing a Lyapunov function and using Foster's theorem.

Given $\mathbf{q} \in \text{int}(C)$, the goal is to show $\mathbf{q}$ can be fulfilled by the MW policy.

By definition of interior there exists $\epsilon > 0$ such that $\mathbf{q}' = \mathbf{q} + \epsilon \mathbf{1} \in C$ where $\mathbf{1} = (1, 1, \cdots, 1)$. We define a Lyapunov function as

$$L(\mathbf{X}(t)) = \sum_{i=1}^{n} X_i(t)^2.$$

In period $t+1$, we have

$$\begin{aligned}
&\mathrm{E}\left[L(\mathbf{X}(t+1)) - L(\mathbf{X}(t))|\mathbf{X}(t)=\mathbf{x}\right] \\
&= \mathrm{E}\left[\sum_{i=1}^{n} X_i(t+1)^2 - X_i(t)^2|\mathbf{X}(t)=\mathbf{x}\right] \\
&\leq \mathrm{E}\left[\sum_{i=1}^{n} (X_i(t)+q_i-V_i(\mathbf{d}(t+1)))^2 - X_i(t)^2|\mathbf{X}(t)=\mathbf{x}\right] \\
&= \mathrm{E}\left[\sum_{i=1}^{n} (q_i-V_i(\mathbf{d}(t+1)))^2 + 2\langle\mathbf{X}(t), \mathbf{q}-\mathbf{V}(\mathbf{d}(t+1))\rangle|\mathbf{X}(t)=\mathbf{x}\right] \\
&\leq \mathrm{E}\left[\sum_{i=1}^{n} (q_i^2+V_i(\mathbf{d}(t+1))^2) + 2\langle\mathbf{X}(t), \mathbf{q}-\mathbf{V}(\mathbf{d}(t+1))\rangle|\mathbf{X}(t)=\mathbf{x}\right] \qquad (10)
\end{aligned}$$

Under the MW policy, by (3) we know that

$$\begin{aligned}
&\mathrm{E}[\langle\mathbf{X}(t), \mathbf{V}(\mathbf{d}(t+1))\rangle|\mathbf{X}(t)=\mathbf{x}] \\
&= \mathrm{E}[\langle\mathbf{X}(t), \mathbf{p}(\mathbf{d}(t+1))\rangle|\mathbf{X}(t)=\mathbf{x}] \\
&= \max_{\mathbf{d}\in D}\langle\mathbf{x}, \mathbf{p}(\mathbf{d})\rangle.
\end{aligned}$$

Since $\mathbf{q}' = \mathbf{q} + \epsilon\mathbf{1} \in C$, we get that

$$\begin{aligned}
&\mathrm{E}[\langle\mathbf{X}(t), \mathbf{q}-\mathbf{V}(\mathbf{d}(t+1))\rangle|\mathbf{X}(t)=\mathbf{x}] \\
&= \mathrm{E}[\langle\mathbf{X}(t), \mathbf{q}'-\mathbf{V}(\mathbf{d}(t+1))\rangle|\mathbf{X}(t)=\mathbf{x}] - \epsilon\langle\mathbf{x}, \mathbf{1}\rangle \\
&= \langle\mathbf{x}, \mathbf{q}'\rangle - \max_{\mathbf{d}\in D}\langle\mathbf{x}, \mathbf{p}(\mathbf{d})\rangle - \epsilon\langle\mathbf{x}, \mathbf{1}\rangle \\
&\leq -\epsilon\langle\mathbf{x}, \mathbf{1}\rangle,
\end{aligned}$$

where the last step is true by Lemma 8.1.

Also since there are finite payoff vectors, we use b_1 to represent an upper bound on all possible payoff values and payoff requirements. Therefore, by (10) we get that

$$\begin{aligned}
\mathrm{E}[L(\mathbf{X}(t+1)) - L(\mathbf{X}(t))|\mathbf{X}(t)=\mathbf{x}] &\leq 2nb_1^2 - 2\epsilon\langle\mathbf{x}, \mathbf{1}\rangle \\
&\leq -1
\end{aligned}$$

for $\mathbf{x}$ satisfying $\langle\mathbf{x}, \mathbf{1}\rangle \geq \frac{nb_1^2}{\epsilon} + \frac{1}{2\epsilon}$.

It is not hard to show[4] there are finite states $\mathbf{x}$ with $\langle\mathbf{x}, \mathbf{1}\rangle < \frac{nb_1^2}{\epsilon} + \frac{1}{2\epsilon}$. Therefore, by Foster's theorem, $\{\mathbf{X}(t)\}_{t\geq 1}$ is positive recurrent and $\mathbf{q}$ is fulfilled by the MW policy. Thus, this shows that $\mathrm{int}(C) \subseteq F_{\mathrm{MW}}$. □

8.3. Proof of Theorem 3.5

We first introduce some further notation. Given two vectors $\boldsymbol{a} = (a_1, a_2, \cdots, a_n)$ and $\mathbf{b} = (b_1, b_2, \cdots, b_n)$, we denote by $\boldsymbol{a} \circ \mathbf{b} = (a_1b_1, a_2b_2, \cdots, a_nb_n)$ the entrywise product.

For any $\mathbf{w} \succ \mathbf{0}$ and $\mathbf{q} \in \mathrm{int}(R_{\mathrm{IB}})$, the goal is to show that $\mathbf{q}$ can be fulfilled by the w-LDF policy.

Let $\boldsymbol{\beta} = (\frac{1}{w_1}, \frac{1}{w_2}, \cdots, \frac{1}{w_n})$. By definition of interior there exists an $\epsilon > 0$ such that $\mathbf{q}' = \mathbf{q} + \epsilon\boldsymbol{\beta} \in R_{\mathrm{IB}}$. By definition of R_{IB}, there exists a vector $\boldsymbol{\alpha} \succ \mathbf{0}$ such that $\mathbf{q}'$ and $\boldsymbol{\alpha}$ satisfy the conditions (4).

[4]This is true because given our assumption that requirement $\mathbf{q}$ and the payoff vectors are rational valued and have finite options, the state space of process $\{\mathbf{X}(t)\}_{t\geq 1}$ is in a lattice, see e.g., [Conway and Sloane 2013].

Consider the following candidate Lyapunov function:

$$L(\mathbf{X}(t)) = \sum_{i=1}^{n} \alpha_i w_i X_i(t)^2.$$

Note that we consider a process $\{\mathbf{X}(t)\}_{t\geq 1}$ that is driven by the w-LDF policy. In period $t+1$, by similar analysis as in (10) we have that

$$\begin{aligned}
&\mathrm{E}\left[L(\mathbf{X}(t+1)) - L(\mathbf{X}(t))|\mathbf{X}(t)=\mathbf{x}\right] \\
&\leq \mathrm{E}\left[\sum_{i=1}^{n} \alpha_i w_i (q_i^2 + V_i(\mathbf{d}(t+1))^2) + 2\langle \boldsymbol{\alpha}\circ\mathbf{w}\circ\mathbf{X}(t), \mathbf{q} - \mathbf{V}(\mathbf{d}(t+1))\rangle | \mathbf{X}(t)=\mathbf{x}\right] \\
&= \mathrm{E}\left[\sum_{i=1}^{n} \alpha_i w_i (q_i^2 + V_i(\mathbf{d}(t+1))^2) + 2\langle \boldsymbol{\alpha}\circ\mathbf{w}\circ\mathbf{X}(t), \mathbf{q}' - \mathbf{V}(\mathbf{d}(t+1))\rangle | \mathbf{X}(t)=\mathbf{x}\right] \\
&\quad -2\epsilon\langle \mathbf{x}, \boldsymbol{\alpha}\rangle \qquad (11)
\end{aligned}$$

Let $\mathbf{d}$ denote the priority decision selected according to w-LDF policy. Thus, We have that

$$\begin{aligned}
&\mathrm{E}\left[\langle \boldsymbol{\alpha}\circ\mathbf{w}\circ\mathbf{X}(t), \mathbf{q}' - \mathbf{V}(\mathbf{d}(t+1))\rangle | \mathbf{X}(t)=\mathbf{x}\right] \\
&= \langle \boldsymbol{\alpha}\circ\mathbf{w}\circ\mathbf{x}, \mathbf{q}' - \mathbf{p}(\mathbf{d})\rangle.
\end{aligned}$$

By reordering users according to priorities, we get

$$\begin{aligned}
&\langle \boldsymbol{\alpha}\circ\mathbf{w}\circ\mathbf{x}, \mathbf{q}' - \mathbf{p}(\mathbf{d})\rangle \\
&= \sum_{i=1}^{n} w_{d_i} x_{d_i} [\alpha_{d_i} q'_{d_i} - \alpha_{d_i} p_{d_i}(\mathbf{d})] \\
&= \sum_{i=1}^{n-1} [w_{d_i} x_{d_i} - w_{d_{i+1}} x_{d_{i+1}}][\sum_{j=1}^{i} \alpha_{d_j} q'_{d_j} - \sum_{j=1}^{i} \alpha_{d_j} p_{d_j}(\mathbf{d})] \\
&\quad + w_{d_n} x_{d_n} [\sum_{j=1}^{n} \alpha_{d_j} q'_{d_j} - \sum_{j=1}^{n} \alpha_{d_j} p_{d_j}(\mathbf{d})].
\end{aligned}$$

By w-LDF policy we know $w_{d_i} x_{d_i} \geq w_{d_{i+1}} x_{d_{i+1}}$. By (4) we have $\sum_{j=1}^{i} \alpha_{d_j} q'_{d_j} \leq \sum_{j=1}^{i} \alpha_{d_j} p_{d_j}(\mathbf{d})$ for $1 \leq i \leq n$. Therefore,

$$\mathrm{E}\left[\langle \boldsymbol{\alpha}\circ\mathbf{w}\circ\mathbf{X}(t), \mathbf{q}' - \mathbf{V}(\mathbf{d}(t+1))\rangle | \mathbf{X}(t)=\mathbf{x}\right] \leq 0.$$

Suppose b_2 is an upper bound on all α_i and w_i, by (11), we get that

$$\begin{aligned}
\mathrm{E}[L(\mathbf{X}(t+1)) - L(\mathbf{X}(t))|\mathbf{X}(t)=\mathbf{x}] \leq & 2nb_2^2 b_1^2 - 2\epsilon\langle \mathbf{x}, \boldsymbol{\alpha}\rangle \\
\leq & -1
\end{aligned}$$

for $\mathbf{x}$ satisfying $\langle \mathbf{x}, \boldsymbol{\alpha}\rangle \geq \frac{nb_2^2 b_1^2}{\epsilon} + \frac{1}{2\epsilon}$.

Again, since there are finite states $\mathbf{x}$ with $\langle \mathbf{x}, \boldsymbol{\alpha}\rangle < \frac{nb_2^2 b_1^2}{\epsilon} + \frac{1}{2\epsilon}$, by Foster's Theorem $\{\mathbf{X}(t)\}_{t\geq 1}$ is positive recurrent and $\mathbf{q}$ is fulfilled by the w-LDF policy.

Therefore, for any $\mathbf{w} \succ \mathbf{0}$, we have that

$$\mathrm{int}(R_{\mathrm{IB}}) \subseteq F_{\text{w-LDF}}.$$

8.4. Proof of Theorem 3.8

The proof of Theorem 3.8 is complicated and we apologize for that. Part of the complication comes from understanding how the property of monotonicity in payoffs characterizes the geometry of the region R_{IB}. Further, a feasible payoff requirement implies feasibilities for all user subsets, and thus we need to look at the projections in all subspaces.

Figure 7 gives the high-level outline for the proof of Theorem 3.8. There are two parts which involve the technical results Lemma 8.2 and 8.3, which we will state in the proof. In order to allow the reader follow the proof, we defer their own proof to later.

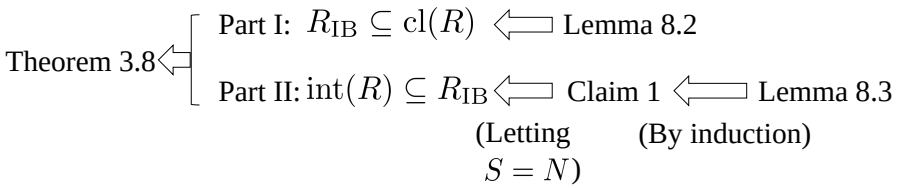
Theorem 3.8
Part I: $R_{\text{IB}} \subseteq \text{cl}(R)$ Lemma 8.2
Part II: $\text{int}(R) \subseteq R_{\text{IB}}$ Claim 1 Lemma 8.3
(Letting $S = N$)
(By induction)

Fig. 7. Outline for the proof of Theorem 3.8.

LEMMA 8.2. *If a system satisfies monotonicity in payoffs, then for all* $\mathbf{q} \in C$ *and all subsets of users* $S \subseteq N$, $\mathbf{q}^S \in C^S$.

We have argued in Section 3.3 that C^S does not necessarily equal to the projection of region C on the subspace of S in general, but the statement is true if the system satisfies monotonicity in payoffs. Please refer to Appendix 8.5 for detailed proof of this lemma.

Given a subset of users $S \subseteq N$, for two vectors $\mathbf{p}^S$ and $\mathbf{q}^S$, we say $\mathbf{p}^S \succ_S \mathbf{q}^S$ if $p_i^S > q_i^S$ for any $i \in S$. We define the subspace of S as $\mathbb{R}_+^S = \{\mathbf{q} \in \mathbb{R}_+^n | q_i = 0, \forall i \notin S\}$.

For a region $X \subseteq \mathbb{R}_+^n$ which lies in the subspace of S, we denote by $\text{int}^S(X)$ the relative interior of X, i.e., the interior of X, relative to the subspace of S. Similarly we use $\text{cl}^S(X)$, $\text{bd}^S(X)$ to represent the relative closure and boundary of X in the subspace of S, respectively.

Following the definition of region R in Definition 3.7, we can express $\text{int}(R)$ and $\text{cl}(R)$ as below,

$$\text{int}(R) = \{\mathbf{q} | \forall S \subseteq \{1, 2, \cdots, n\}, \mathbf{q}^S \in \text{int}^S(C^S \setminus B^S)\}, \tag{12}$$

$$\text{cl}(R) = \{\mathbf{q} | \forall S \subseteq \{1, 2, \cdots, n\}, \mathbf{q}^S \in \text{cl}^S(C^S \setminus B^S)\}. \tag{13}$$

By definition, we know C^S and B^S are closed sets. We can also show[5]

$$\text{int}^S(C^S \setminus B^S) = \text{int}^S(C^S) \setminus B^S,$$
$$\text{cl}^S(C^S \setminus B^S) = C^S \setminus \text{int}^S(B^S).$$

Next we prove Theorem 3.8 in two parts: $R_{\text{IB}} \subseteq \text{cl}(R)$ and $\text{int}(R) \subseteq R_{\text{IB}}$.

Part I of the proof:

We start with the easy part and first show $R_{\text{IB}} \subseteq \text{cl}(R)$.

[5]This is clear from the definition of C^S and B^S. We omit the proof to save space.

Given $\mathbf{q} \in R_{\text{IB}}$, by definition, there exists $\boldsymbol{\alpha} \succ \mathbf{0}$ such that for all subsets of users S,

$$\sum_{i \in S} \alpha_i q_i \leq \min_{\mathbf{d} \in D(S)} \sum_{i \in S} \alpha_i p_i(\mathbf{d}),$$

i.e.,

$$\langle \boldsymbol{\alpha}^S, \mathbf{q}^S \rangle \leq \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle. \tag{14}$$

To show $\mathbf{q} \in \text{cl}(R)$, by (13) we need to show for all subsets of users S that $\mathbf{q}^S \in C^S \setminus \text{int}^S(B^S)$.

Since $\mathbf{q} \in R_{\text{IB}} \subseteq C$, by Lemma 8.2 we have $\mathbf{q}^S \in C^S$.

Now suppose for some subset of users S, $\mathbf{q}^S \in \text{int}^S(B^S)$. By definition of B^S and interior, there exists $\mathbf{x}^S \in \text{Conv}(P^S)$ such that $\mathbf{q}^S \succeq \mathbf{x}^S$ and $\mathbf{q}^S \neq \mathbf{x}^S$, which implies that

$$\langle \boldsymbol{\alpha}^S, \mathbf{q}^S \rangle > \langle \boldsymbol{\alpha}^S, \mathbf{x}^S \rangle.$$

Since $\mathbf{x}^S \in \text{Conv}(P^S)$, there exists $\{c_{\mathbf{d}} | \mathbf{d} \in D(S)\}$ such that $\mathbf{x}^S = \sum_{\mathbf{d} \in D(S)} c_{\mathbf{d}} \mathbf{p}^S(\mathbf{d})$ and $c_{\mathbf{d}} \geq 0$, $\sum_{\mathbf{d} \in D(S)} c_{\mathbf{d}} = 1$. Therefore,

$$\begin{aligned} \langle \boldsymbol{\alpha}^S, \mathbf{q}^S \rangle > \langle \boldsymbol{\alpha}^S, \mathbf{x}^S \rangle &= \sum_{\mathbf{d} \in D(S)} c_{\mathbf{d}} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle \\ &\geq \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle, \end{aligned} \tag{15}$$

which is contradicted to (14). Therefore, $\mathbf{q}^S \notin \text{int}^S(B^S)$ and thus $\mathbf{q} \in \text{cl}(R)$, implying that $R_{\text{IB}} \subseteq \text{cl}(R)$.

Part II of the proof:

For the second part we show $\text{int}(R) \subseteq R_{\text{IB}}$.

Given $\mathbf{q} \in \text{int}(R)$, by (12), we know for all S, $\mathbf{q}^S \in \text{int}^S(C^S) \setminus B^S$. The goal is to show $\mathbf{q} \in R_{\text{IB}}$, i.e., to find an $\boldsymbol{\alpha} \succ \mathbf{0}$ such that for all subsets of users S,

$$\langle \boldsymbol{\alpha}^S, \mathbf{q}^S \rangle \leq \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle. \tag{16}$$

We start with the following lemma which is proved in Appendix 8.6 in the sequel.

LEMMA 8.3. *If a system satisfies monotonicity in payoffs, given $\mathbf{q} \in int(R)$, for each subset of users S, there exists nonzero $\boldsymbol{\beta}^S \succeq \mathbf{0}$, such that for all $S' \subseteq S$ where $\boldsymbol{\beta}^{S'} \neq \mathbf{0}$,*

$$\langle \boldsymbol{\beta}^{S'}, \mathbf{q}^{S'} \rangle < \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\beta}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle. \tag{17}$$

Given Lemma 8.3, by letting $S = N$ we find $\boldsymbol{\beta}$ that is very similar to the $\boldsymbol{\alpha}$ we are looking for, except for two differences: (1) $\boldsymbol{\beta}$ may not be strictly positive, and (2) it is strictly "less than" in (17). The idea is to add a small perturbation to $\boldsymbol{\beta}$ to construct a strictly positive vector.

Formally, given Lemma 8.3, to show $\mathbf{q} \in R_{\text{IB}}$ we shall prove the following even stronger statement by induction.

Claim 1: If a system satisfies monotonicity in payoffs, given $\mathbf{q} \in \text{int}(R)$, for each subset of users S, there exists $\boldsymbol{\alpha}^S \succ_S \mathbf{0}$ such that for all $S' \subseteq S$,

$$\langle \boldsymbol{\alpha}^{S'}, \mathbf{q}^{S'} \rangle \leq \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\alpha}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle. \tag{18}$$

By Claim 1, let $S = N$, we can find $\boldsymbol{\alpha} \succ \mathbf{0}$ satisfying (16) which implies $\mathbf{q} \in R_{\text{IB}}$. Therefore, it suffices to prove Claim 1. We prove this by induction on the cardinality $|S|$ of user set S.

If $|S| = 0$, clearly Claim 1 is correct.

Suppose Claim 1 is correct for all S with $|S| \leq k-1$ where $k \geq 1$. Given an S with $|S| = k$, by Lemma 8.3, we can find nonzero $\boldsymbol{\beta}^S \succeq \mathbf{0}$ satisfying the conditions in Lemma 8.3. We separate the set S into two sets S_1 and S_2 where

$$\beta_i^S > 0, i \in S_1,$$
$$\beta_i^S = 0, i \in S_2.$$

Since $\boldsymbol{\beta}^S \neq \mathbf{0}$, $|S_2| \leq |S| - 1 = k - 1$. By induction of Claim 1 on S_2, there exists $\boldsymbol{\gamma}^{S_2} \succ_{S_2} \mathbf{0}$ such that for any $S' \subseteq S_2$,

$$\langle \boldsymbol{\gamma}^{S'}, \mathbf{q}^{S'} \rangle \leq \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\gamma}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle. \tag{19}$$

We claim that for small enough $\delta > 0$,

$$\boldsymbol{\alpha}^S = \boldsymbol{\beta}^S + \delta \boldsymbol{\gamma}^{S_2} \succ_S \mathbf{0}$$

satisfies condition (18) for all $S' \subseteq S$.

Any $S' \subseteq S$ falls into one of the following two cases: $S' \subseteq S_2$ and $S' \not\subseteq S_2$. It suffices to show (18) in each case.

If $S' \subseteq S_2$, then $\boldsymbol{\alpha}^{S'} = \boldsymbol{\gamma}^{S'}$. By (19), we know (18) is correct.

If $S' \not\subseteq S_2$, then $\boldsymbol{\beta}^{S'} \neq \mathbf{0}$. Let $\boldsymbol{\gamma}^{S'} = (\boldsymbol{\gamma}^{S_2})^{S'}$. We know

$$\langle \boldsymbol{\alpha}^{S'}, \mathbf{q}^{S'} \rangle = \langle \boldsymbol{\beta}^{S'}, \mathbf{q}^{S'} \rangle + \delta \langle \boldsymbol{\gamma}^{S'}, \mathbf{q}^{S'} \rangle.$$

By Lemma 8.3, $\langle \boldsymbol{\beta}^{S'}, \mathbf{q}^{S'} \rangle < \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\beta}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle$. Since there are finite subsets S', for small enough δ,

$$\begin{aligned} \langle \boldsymbol{\alpha}^{S'}, \mathbf{q}^{S'} \rangle &\leq \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\beta}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle \\ &\leq \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\alpha}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle, \end{aligned}$$

i.e., (18) holds true.

In summary, this proves Claim 1 and thus $\mathbf{q} \in R_{\text{IB}}$. Therefore, $\text{int}(R) \subseteq R_{\text{IB}}$.

8.5. Proof of Lemma 8.2

First we introduce a further notation. Given a decision $\mathbf{d}$ and a user set S, we let $m(\mathbf{d}, S)$ represent the decision that satisfies

— $m(\mathbf{d}, S) \in D(S)$.

— For users $i, j \in S$ or $i, j \notin S$, if i has higher priority than j in $\mathbf{d}$, then i also has higher priority than j in decision $m(\mathbf{d}, S)$.

In other words, $m(\mathbf{d}, S)$ is the priority decision obtained by modifying decision $\mathbf{d}$ to assign highest priorities to users in S without changing the relative orders in and out of S, respectively.

Given that the system satisfies monotonicity in payoffs, for all $i \in S$, we have that

$$p_i(m(\mathbf{d}, S)) \geq p_i(\mathbf{d}). \tag{20}$$

Given $\mathbf{q} \in C$, the goal is to show $\mathbf{q}^S \in C^S$ for all subsets of users S. By definition of C, there exists a convex combination of vectors in P that dominates $\mathbf{q}$, i.e., there exists $\{\alpha_{\mathbf{d}} | \mathbf{d} \in D\}$ such that

$$\mathbf{q} \preceq \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} \mathbf{p}(\mathbf{d}),$$

and $\alpha_{\mathbf{d}} \geq 0, \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} = 1$.

Therefore, for any subset of users S,

$$\mathbf{q}^S \preceq \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} \mathbf{p}^S(\mathbf{d}),$$

which by (20) gives

$$\mathbf{q}^S \preceq \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} \mathbf{p}^S(m(\mathbf{d}, S)).$$

We let $\mathbf{x}^S = \sum_{\mathbf{d} \in D} \alpha_{\mathbf{d}} \mathbf{p}^S(m(\mathbf{d}, S))$. Since $m(\mathbf{d}, S) \in D(S)$, we know $\mathbf{x}^S \in \text{Conv}(P^S)$, and by definition of C^S,

$$\mathbf{q}^S \in C^S.$$

8.6. Proof of Lemma 8.3

Given $\mathbf{q} \in \text{int}(R)$, by (12) we know that for all subsets of users S, $\mathbf{q}^S \in \text{int}^S(C^S) \setminus B^S$. Given a subset of users S, the goal is to find $\boldsymbol{\beta}^S$ which satisfies the requirements in Lemma 8.3. In this proof, we focus on the subspace of S.

Since $\mathbf{q}^S \in \text{int}^S(C^S)$, there exists $\mathbf{x}^S \in \text{Conv}(P^S)$ such that $\mathbf{q}^S \prec_S \mathbf{x}^S$. Since $\mathbf{q}^S \notin B^S$ and $\mathbf{x}^S \in \text{Conv}(P^S) \subseteq B^S$, we know that connecting $\mathbf{q}^S$ and $\mathbf{x}^S$ intersects $\text{bd}^S(B^S)$ at some point denoted by $\mathbf{v}^S$. By the closure property of B^S, $\mathbf{v}^S \in B^S$ and thus, $\mathbf{v}^S \succ_S \mathbf{q}^S$. Since $\mathbf{v}^S \in \text{bd}^S(B^S)$, we get that $\mathbf{v}^S$ lies on a supporting hyperplane [Boyd and Vandenberghe 2009] of B^S, and by definition of B^S, there exists nonzero normal vector $\boldsymbol{\beta}^S \succeq \mathbf{0}$ of this supporting hyperplane such that

$$\langle \boldsymbol{\beta}^S, \mathbf{v}^S \rangle = \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}) \rangle. \tag{21}$$

Figure 8 conceptually shows the process of constructing $\boldsymbol{\beta}^S$. The circles represent the expected payoff vectors. For simplicity we suppress the superscript S in the figure. We shall show this $\boldsymbol{\beta}^S$ satisfies the requirements in Lemma 8.3.

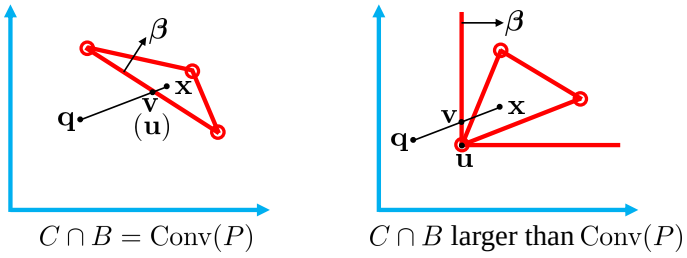

Fig. 8. The process of constructing $\boldsymbol{\beta}$ when $C \cap B$ equals to $\text{Conv}(P)$ (left figure), and when $C \cap B$ is larger than $\text{Conv}(P)$ (right figure).

Since $\mathbf{v}^S \in \text{bd}^S(B^S) \subseteq B^S$, there exists $\mathbf{u}^S \in \text{Conv}(P^S)$ such that $\mathbf{v}^S \succeq \mathbf{u}^S$, and by (21), we get that

$$\langle \boldsymbol{\beta}^S, \mathbf{u}^S \rangle \leq \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}) \rangle.$$

The vector $\mathbf{u}^S$ is also shown in Figure 8.

On the other hand, since $\mathbf{u}^S \in \text{Conv}(P^S)$, by similar analysis as in (15) we know that

$$\langle \boldsymbol{\beta}^S, \mathbf{u}^S \rangle \geq \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}) \rangle.$$

Thus,

$$\langle \boldsymbol{\beta}^S, \mathbf{u}^S \rangle = \langle \boldsymbol{\beta}^S, \mathbf{v}^S \rangle = \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}) \rangle,$$

which implies $u_i^S = v_i^S$ if $\beta_i^S \neq 0$. Since $\mathbf{v}^S \succ_S \mathbf{q}^S$, for all subset $S' \subseteq S$ where $\boldsymbol{\beta}^{S'} \neq \mathbf{0}$, we have that

$$\langle \boldsymbol{\beta}^{S'}, \mathbf{q}^{S'} \rangle < \langle \boldsymbol{\beta}^{S'}, \mathbf{u}^{S'} \rangle,$$

Therefore, to show (17) it suffices to show for all $S' \subseteq S$ where $\boldsymbol{\beta}^{S'} \neq \mathbf{0}$ that

$$\langle \boldsymbol{\beta}^{S'}, \mathbf{u}^{S'} \rangle \leq \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\beta}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle. \tag{22}$$

Given $\mathbf{u}^S \in \text{Conv}(P^S)$, we write $\mathbf{u}^S = \sum_{\mathbf{d} \in D(S)} c_{\mathbf{d}} \mathbf{p}^S(\mathbf{d})$ where $c_{\mathbf{d}} \geq 0$ and $\sum_{\mathbf{d} \in D(S)} c_{\mathbf{d}} = 1$.

For all $S' \subseteq S$ where $\boldsymbol{\beta}^{S'} \neq \mathbf{0}$, we can rewrite $\langle \boldsymbol{\beta}^{S'}, \mathbf{u}^{S'} \rangle$ as follows,

$$\begin{aligned}
\langle \boldsymbol{\beta}^{S'}, \mathbf{u}^{S'} \rangle &= \langle \boldsymbol{\beta}^S, \mathbf{u}^S \rangle - \langle \boldsymbol{\beta}^{S \setminus S'}, \mathbf{u}^{S \setminus S'} \rangle \\
&= \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}) \rangle - \sum_{\mathbf{d} \in D(S)} c_{\mathbf{d}} \langle \boldsymbol{\beta}^{S \setminus S'}, \mathbf{p}^{S \setminus S'}(\mathbf{d}) \rangle \\
&\leq \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}) \rangle - \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^{S \setminus S'}, \mathbf{p}^{S \setminus S'}(\mathbf{d}) \rangle.
\end{aligned}$$

To show (22), it suffices to show that

$$\begin{aligned}
&\min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}) \rangle \\
&\leq \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^{S \setminus S'}, \mathbf{p}^{S \setminus S'}(\mathbf{d}) \rangle &(23)\\
&\quad + \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\beta}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle. &(24)
\end{aligned}$$

Suppose $\mathbf{d}_1 \in D(S)$ and $\mathbf{d}_2 \in D(S')$ are the optimal solutions for (23) and (24), respectively. Formally,

$$\langle \boldsymbol{\beta}^{S \setminus S'}, \mathbf{p}^{S \setminus S'}(\mathbf{d}_1) \rangle = \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\beta}^{S \setminus S'}, \mathbf{p}^{S \setminus S'}(\mathbf{d}) \rangle,$$

$$\langle \boldsymbol{\beta}^{S'}, \mathbf{p}^{S'}(\mathbf{d}_2) \rangle = \min_{\mathbf{d} \in D(S')} \langle \boldsymbol{\beta}^{S'}, \mathbf{p}^{S'}(\mathbf{d}) \rangle.$$

We consider the unique decision $\mathbf{d}_3$ that satisfies the following: First, $\mathbf{d}_3 \in D(S')$, i.e., $\mathbf{d}_3$ assigns highest priority to users in S'. Second, the priority ordering for user subset S' in $\mathbf{d}_3$ are the same as those in $\mathbf{d}_2$. Third, the priority ordering for user subset $N \setminus S'$ in $\mathbf{d}_3$ are the same as those in $\mathbf{d}_1$.

Since $\mathbf{d}_1 \in D(S)$, we know $\mathbf{d}_3 \in D(S)$ and therefore,

$$\min_{\mathbf{d}\in D(S)} \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d})\rangle \le \langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}_3)\rangle.$$

Now it suffices to show that

$$\langle \boldsymbol{\beta}^S, \mathbf{p}^S(\mathbf{d}_3)\rangle \le \langle \boldsymbol{\beta}^{S\setminus S'}, \mathbf{p}^{S\setminus S'}(\mathbf{d}_1)\rangle + \langle \boldsymbol{\beta}^{S'}, \mathbf{p}^{S'}(\mathbf{d}_2)\rangle.$$

This is true because given that the system satisfies monotonicity in payoffs, we can get that

$$p_i^{S'}(\mathbf{d}_3) \le p_i^{S'}(\mathbf{d}_2) \text{ for } i \in S',$$

and

$$p_i^{S\setminus S'}(\mathbf{d}_3) \le p_i^{S\setminus S'}(\mathbf{d}_1) \text{ for } i \in S \setminus S'.$$

In summary, this proves (22) and thus $\boldsymbol{\beta}^S$ satisfies the conditions in Lemma 8.3.

8.7. Proof of Theorem 3.10

Clearly $F_{\text{w-LDF}} \subseteq F \subseteq \text{cl}(C)$. To show $\text{int}(C) \subseteq F_{\text{w-LDF}}$, by (5) it suffices to show $\text{int}(C) \subseteq \text{int}(R)$.

Given $\mathbf{q} \in \text{int}(C)$, the goal is to show $\mathbf{q} \in \text{int}(R)$, i.e., for all user subsets S,

$$\mathbf{q}^S \in \text{int}^S(C^S) \setminus B^S.$$

Given a user subset S, by Lemma 8.2 we know $\mathbf{q}^S \in C^S$. Further we can show $\mathbf{q}^S \in \text{int}^S(C^S)$ since otherwise $\mathbf{q} \in \text{bd}(C)$. By definition of interior and C^S, there exists $\mathbf{x}^S \in \text{Conv}(P^S)$ such that $\mathbf{q}^S \prec_S \mathbf{x}^S$.

Suppose $\mathbf{q}^S \in B^S$, by definition there exists $\mathbf{y}^S \in \text{Conv}(P^S)$ such that $\mathbf{q}^S \succeq \mathbf{y}^S$. Now we get two vectors $\mathbf{x}^S, \mathbf{y}^S \in \text{Conv}(P^S)$ and $\mathbf{x}^S \succ_S \mathbf{y}^S$. Since vectors in P^S lie on a hyperplane and $\mathbf{x}^S, \mathbf{y}^S \in \text{Conv}(P^S)$, there exists nonzero $\boldsymbol{\alpha}^S \succeq \mathbf{0}$ such that

$$\langle \boldsymbol{\alpha}^S, \mathbf{x}^S\rangle = \langle \boldsymbol{\alpha}^S, \mathbf{y}^S\rangle,$$

which contradicts with $\mathbf{x}^S \succ_S \mathbf{y}^S$.

Therefore, $\mathbf{q}^S \notin B^S$ and thus $\text{int}(C) \subseteq \text{int}(R)$.

8.8. Proof of Theorem 3.14

Given monotonicity in payoffs, by (5) we know $\text{int}(R) \subseteq F_{\text{w-LDF}}$ and F differs from C by at most a boundary. By the definition of the efficiency ratio, we know

$$\gamma_{\text{w-LDF}} = \sup\{\gamma|\gamma F \subseteq F_{\text{w-LDF}}\} \ge \sup\{\gamma|\gamma C \subseteq \text{int}(R)\} = \sup\{\gamma|\gamma C \subseteq \text{cl}(R)\}.$$

To show $\gamma_{\text{w-LDF}} \ge \min_{S\subseteq N} \sigma_S$, it suffices to show $\sup\{\gamma|\gamma C \subseteq \text{cl}(R)\} \ge \min_{S\subseteq N} \sigma_S$, i.e., for each $\mathbf{q} \in C$, we have $\min_{S\subseteq N} \sigma_S \cdot \mathbf{q} \in \text{cl}(R)$, which is equivalent to showing that for each $\mathbf{q} \in C$, there exists a subset of users $S \subseteq N$, such that $\sigma_S \cdot \mathbf{q} \in \text{cl}(R)$.

Given a $\mathbf{q} \in C$, we define $\lambda(\mathbf{q}, R) = \sup\{\lambda|\lambda\mathbf{q} \in R\}$ which represents how far the vector $\mathbf{q}$ can extend before it goes beyond the region R and let $\mathbf{q}_{(R)} = \lambda(\mathbf{q}, R) \cdot \mathbf{q}$. We claim there exists a user subset S such that $\mathbf{q}_{(R)}^S \in \text{bd}^S(B^S)$ since otherwise we can increase $\lambda(\mathbf{q}, R)$ while guaranteeing $\mathbf{q}_{(R)}$ is still in R. Next we shall show $\sigma_S \cdot \mathbf{q} \in \text{cl}(R)$, i.e., $\lambda(\mathbf{q}, R) \ge \sigma_S$.

Since $\mathbf{q}_{(R)}^S \in \text{bd}^S(B^S) \subseteq B^S$, there exists $\mathbf{x}^S \in \text{Conv}(P^S)$ such that $\mathbf{q}_{(R)}^S \succeq \mathbf{x}^S$. Since $\mathbf{q} \in C$, by Lemma 8.2 we know $\mathbf{q}^S \in C^S$ and thus, there exists $\mathbf{y}^S \in \text{Conv}(P^S)$ such that $\mathbf{q}^S \preceq \mathbf{y}^S$.

Given that $\mathbf{q}_{(R)} = \lambda(\mathbf{q}, R) \cdot \mathbf{q}$, for any nonzero $\boldsymbol{\alpha}^S \succeq \mathbf{0}$, we have that

$$\langle \mathbf{q}^S_{(R)}, \boldsymbol{\alpha}^S \rangle = \lambda(\mathbf{q}, R) \cdot \langle \mathbf{q}^S, \boldsymbol{\alpha}^S \rangle.$$

By $\mathbf{q}^S_{(R)} \succeq \mathbf{x}^S$ and $\mathbf{q}^S \preceq \mathbf{y}^S$, we get that

$$\langle \mathbf{x}^S, \boldsymbol{\alpha}^S \rangle \leq \lambda(\mathbf{q}, R) \cdot \langle \mathbf{y}^S, \boldsymbol{\alpha}^S \rangle.$$

Since $\mathbf{x}^S \in \text{Conv}(P^S)$, by similar analysis as in (15) we know $\langle \mathbf{x}^S, \boldsymbol{\alpha}^S \rangle \geq \min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle$. Similarly we can show $\langle \mathbf{y}^S, \boldsymbol{\alpha}^S \rangle \leq \max_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle$. Thus,

$$\min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle \leq \lambda(\mathbf{q}, R) \cdot \max_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle. \tag{25}$$

Clearly, $\max_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle > 0$ for any nonzero $\boldsymbol{\alpha}^S \succeq \mathbf{0}$.

Since (25) is true for any nonzero $\boldsymbol{\alpha}^S \succeq \mathbf{0}$, we get that

$$\lambda(\mathbf{q}, R) \geq \max_{\substack{\boldsymbol{\alpha}^S \succeq \mathbf{0} \\ \boldsymbol{\alpha}^S \neq \mathbf{0}}} \frac{\min_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle}{\max_{\mathbf{d} \in D(S)} \langle \boldsymbol{\alpha}^S, \mathbf{p}^S(\mathbf{d}) \rangle} = \sigma_S.$$

Therefore, for each $\mathbf{q} \in C$, there exists a subset of users $S \subseteq N$ such that $\sigma_S \cdot \mathbf{q} \in \text{cl}(R)$, and thus, $\gamma_{\text{w-LDF}} \geq \min_{S \subseteq N} \sigma_S$.

8.9. Proof of Corollary 4.2

By Theorem 3.10, to show w-LDF policies are feasibility optimal, it suffices to show the system satisfies subset payoff equivalence. To show this, it suffices to show for all user subsets $S \subseteq N$ and all priority decisions $\mathbf{d}_1, \mathbf{d}_2 \in D(S)$ that

$$\langle \mathbf{1}^S, \mathbf{p}^S(\mathbf{d}_1) \rangle = \langle \mathbf{1}^S, \mathbf{p}^S(\mathbf{d}_2) \rangle.$$

This is true because we can convert $\mathbf{d}_1$ to $\mathbf{d}_2$ by repeatedly switching a pair of users in $\mathbf{d}_1$ at each step such that at step k both decisions assign the highest k priorities to the same users, respectively. By the definition of exchangeable expected payoffs, the sum of the expected payoffs for users in S remains the same at each step.